Two-message quantum interactive proofs and the quantum separability problem

Patrick Hayden, Kevin Milner, and Mark M. Wilde
School of Computer Science,
McGill University,
Montreal, Quebec H3A 2A7, Canada

December 13, 2012

Abstract

Suppose that a polynomial-time mixed-state quantum circuit, described as a sequence of local unitary interactions followed by a partial trace, generates a quantum state shared between two parties. One might then wonder, does this quantum circuit produce a state that is separable or entangled? Here, we give evidence that it is computationally hard to decide the answer to this question, even if one has access to the power of quantum computation. We begin by exhibiting a two-message quantum interactive proof system that can decide the answer to a promise version of the question. We then prove that the promise problem is hard for the class of promise problems with "quantum statistical zero knowledge" (QSZK) proof systems by demonstrating a polynomial-time Karp reduction from the QSZK-complete promise problem "quantum state distinguishability" to our quantum separability problem. By exploiting Knill's efficient encoding of a matrix description of a state into a description of a circuit to generate the state, we can show that our promise problem is NP-hard. Thus, the quantum separability problem (as phrased above) constitutes the first nontrivial promise problem decidable by a two-message quantum interactive proof system while being hard for both NP and QSZK. We consider a variant of the problem, in which a given polynomial-time mixed-state quantum circuit accepts a quantum state as input, and the question is to decide if there is an input to this circuit which makes its output separable across some bipartite cut. We prove that this problem is a complete promise problem for the class QIP of problems decidable by quantum interactive proof systems. Finally, we show that a two-message quantum interactive proof system can also decide a multipartite generalization of the quantum separability problem.

1 Introduction

Quantum entanglement plays a central role in quantum information science [HHHH09]. It is believed to be one of the reasons behind the computational power of quantum algorithms [EJ98], it can enhance the communication capacities of channels in fascinating ways [BSST99, BSST02, CLMW10, PLM$^+$11], most notably as exhibited in quantum teleportation [BBC$^+$93] and super-dense coding [BW92], and it is a resource for device-independent quantum key distribution [Eke91, VV12]. For these reasons and others, the characterization and systematic understanding of entanglement have long been important goals of quantum information and complexity theory.

Every quantum state has a mathematical description as a density operator ρ, which is a unit-trace, positive semidefinite operator acting on some Hilbert space $\mathcal{H}$. (In this work, we restrict

ourselves to finite-dimensional Hilbert spaces.) If the Hilbert space $\mathcal{H}$ has a factorization as a tensor product $\mathcal{H}_A \otimes \mathcal{H}_B$ of two Hilbert spaces $\mathcal{H}_A$ and $\mathcal{H}_B$, then we write the state ρ as ρ_{AB} and say that it is *separable* if it admits a decomposition of the following form:

$$\rho_{AB} = \sum_{z \in \mathcal{Z}} p_Z(z) \ \sigma_A^z \otimes \tau_B^z, \tag{1}$$

for collections $\{\sigma_A^z\}$ and $\{\tau_B^z\}$ of quantum states and some probability distribution $p_Z(z)$ over an alphabet $\mathcal{Z}$ [Wer89]. Let $\mathcal{S}$ denote the set of separable states. Often, we think of the systems A and B (corresponding to Hilbert spaces $\mathcal{H}_A$ and $\mathcal{H}_B$) as being spatially separated, with an experimentalist Alice possessing system A, while her colleague Bob possesses system B. The intuition behind the above definition of separability is that a separable state can be prepared without any quantum interaction between systems A and B. That is, there is an effectively classical procedure by which they can prepare a separable state: Alice selects a classical variable z according to $p_Z(z)$, prepares the state σ^z in her lab, and sends Bob the variable z so that he can prepare the state τ^z in his lab. After this process, they both discard the variable z, so that (1) describes their shared quantum state. It is clear from inspecting (1) that the set of separable states is a convex set, and that any separable state has the following form as well:

$$\rho_{AB} = \sum_{x \in \mathcal{X}} p_X(x) \ |\psi_x\rangle \langle\psi_x|_A \otimes |\phi_x\rangle \langle\phi_x|_B, \tag{2}$$

for collections $\{|\psi_x\rangle_A\}$ and $\{|\phi_x\rangle_B\}$ of pure states and some probability distribution $p_X(x)$ over an alphabet $\mathcal{X}$ of size no larger than $\dim(\mathcal{H})^2$, the cardinality bound following from Caratheodory's theorem [Hor97]. Finally, if the state ρ_{AB} does not admit a decomposition of the above form, then it is *entangled*—it can only be prepared by a quantum interaction between systems A and B.

Given the many applications of entanglement, it is clearly important to be able to decide if a particular bipartite state is separable or entangled. When the state is specified as the rational entries of a density matrix acting on a finite-dimensional Hilbert space $\mathcal{H} = \mathcal{H}_A \otimes \mathcal{H}_B$, one can formulate several variations of the problem, all of them being known collectively as the *quantum separability problem*, and characterize their computational complexity [Gur03, Gha10, BCY11b]. (See Ref. [Ioa07] for a useful, though now somewhat outdated review.) Gurvits proved that it is NP-hard to decide if a state $\rho_{AB} \in \mathcal{S}$ or if

$$\min_{\sigma_{AB} \in \mathcal{S}} \|\rho_{AB} - \sigma_{AB}\|_2 \geq \varepsilon,$$

where $\|A\|_2 \equiv \sqrt{\text{Tr}\{A^\dagger A\}}$ is the Hilbert-Schmidt norm and ε is some positive number larger than an inverse exponential in $\dim(\mathcal{H})$. Gharibian later improved upon this result by showing that this formulation of the quantum separability problem is strongly NP-hard—it is still NP-hard even if ε is promised to be larger than an inverse polynomial in $\dim(\mathcal{H})$. Brandão, Christandl, and Yard then offered a quasi-polynomial time algorithm that decides the quantum separability problem if it is promised that ε is a positive constant [BCY11b], by appealing to their Pinsker-inequality-like lower bound on the squashed entanglement [BCY11a] and to the k-extendibility separability test of Doherty *et al.* [DPS02, DPS04]. (These concepts are defined later in our paper.) They also considered a variant of the promise problem where the Hilbert-Schmidt distance is replaced by the one-way LOCC distance [MWW09], which characterizes the distinguishability of ρ_{AB} and $\mathcal{S}$ if Alice and Bob are allowed to perform local operations and to send one message of classical communication (from either Alice to Bob or Bob to Alice).

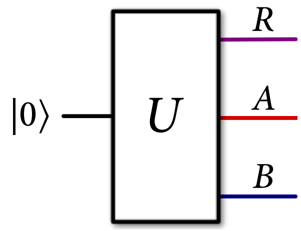

Figure 1: A unitary circuit which generates a bipartite state ρ_{AB} on systems A and B. The qubits in the reference system R are traced out.

In the circuit model of quantum computation, quantum states are generated by unitary circuits acting on some number of qubits (with some of them being traced out in the mixed-state circuit model [AKN98]), and we measure the complexity of a quantum computation by how the circuit size (number of gates and wires) scales with the length of the input [Wat09a]. (Note that if the circuit size is polynomial in the input length, then the number of qubits on which the circuit acts is likewise polynomial in the input length.) Thus, from the perspective of quantum computational complexity theory [Wat09a], one might consider the prior formulations of the quantum separability problem to be somewhat restrictive. The reason is the same as that given in [Ros09]: the mathematical description of a bipartite quantum state is polynomial in the dimension of the Hilbert space, but this Hilbert space is exponential in the number of qubits in the state. Thus, the matrix representation is exponentially larger than it needs to be when we are in the setting of the circuit model of quantum computation. Also, the circuit model is natural physically, as the evolution induced by a time-varying two-body Hamiltonian can be efficiently described by a quantum circuit [BACS07].

With this dual computational and physical motivation in mind, we take an approach to the quantum separability problem along the above lines. We suppose that we are given a description of a quantum circuit of polynomial size as a sequence of quantum gates chosen from some finite gate set, and the circuit acts on a number of qubits, each in the state $|0\rangle$. As part of the description, the qubits are divided into three sets: a set of reference system qubits which are traced out, a set of qubits which belongs to Alice, and another set which belongs to Bob. Figure 1 depicts such a circuit. We also impose the following promise regarding the state ρ_{AB} generated by the circuit: there is either some state in the set of separable states which is ε_1-close to ρ_{AB} in the trace distance:

$$\min_{\sigma_{AB}\in\mathcal{S}} \|\rho_{AB} - \sigma_{AB}\|_1 \leq \varepsilon_1, \tag{3}$$

or every state in this set is at least ε_2-away from ρ_{AB} in the one-way LOCC distance:

$$\min_{\sigma_{AB}\in\mathcal{S}} \|\rho_{AB} - \sigma_{AB}\|_{1-\mathrm{LOCC}} \geq \varepsilon_2, \tag{4}$$

and the difference $\varepsilon_2 - \varepsilon_1$ is larger than a sufficiently large inverse polynomial in the circuit size. The reason for the asymmetric choice of norms will become clear later in the article, but for the moment we will simply note that the problem is well-defined: the two sets identified by (3) and (4) do not overlap because $\|*\|_{1-\mathrm{LOCC}} \leq \|*\|_1$ [MWW09]. Let $\mathsf{QSEP\text{-}CIRCUIT}(\varepsilon_1, \varepsilon_2)$ denote this promise problem, and we will also refer to it throughout this paper as both QSEP-CIRCUIT and the quantum separability problem.

Quantum interactive proof systems [Wat03, KW00, Wat09a] provide a natural framework for analyzing QSEP-CIRCUIT. The idea is that a *verifier*, who has access to a polynomial-time bounded quantum computer, can exchange quantum messages with a *prover*, who has access to a computationally unbounded quantum computer, in order to be convinced by the prover of the truth of some statement. Crucial in this setting are the notions of *interaction* and *verification*: in the case of a positive instance of a problem, there exists some strategy by which the prover can convince the verifier to accept with high probability, whereas in the case of a negative instance, with high probability, there is nothing that the prover can do to convince the verifier to accept. One appeal of the quantum setting is that it is in some ways conceptually simpler than its classical counterpart, thanks to Kitaev and Watrous' demonstration that it is possible to parallelize any quantum interactive proof system to at most three messages between the verifier and prover [KW00]. Their result leaves just four natural complexity classes derived from quantum interactive proof systems with a single prover:

- QIP(3)=QIP—the full power of quantum interactive proof systems (known to be equivalent to PSPACE [JJUW10]),
- QIP(2)—there are two messages exchanged (verifier-to-prover followed by prover-to-verifier),
- QIP(1)—the prover sends a quantum state to the verifier (this class is more commonly known as QMA),
- QIP(0)—the verifier proceeds on his own (this is equivalent to BQP, the class of problems efficiently decidable on a quantum computer with bounded error).

It has been remarked that QIP(2) appears to be the "most mysterious" of these complexity classes [JUW09]. At the very least, not much is known about it.

2 Overview of Results

In this paper, we contribute the following results:

- QSEP-CIRCUIT is decidable by a two-message quantum interactive proof system for a wide range of parameters. The proof system builds upon the approach of Brandão *et al.* [BCY11b] and the notion of k-extendibility [DPS02, DPS04]. In particular, the verifier generates the state ρ_{AB} using its description as a quantum circuit and then sends the reference system to the prover. In the case of a positive instance, the state is separable (or close to it), and the prover can generate a purification of a k-extension of the state by acting with a unitary operation on the reference system and some ancilla qubits. The prover sends all of his output qubits back to the verifier, who then performs phase estimation over the symmetric group [Kit95] (also known as the "permutation test" [BBD$^+$97, KNY08]) in order to verify whether the state sent by the prover is a k-extension. This proof system has completeness and soundness error directly related to ε_1 in (3) and ε_2 in (4).
- QSEP-CIRCUIT is hard for the class of promise problems decidable by a quantum statistical zero-knowledge (QSZK) proof system [Wat02, Wat09b]. A QSZK proof system is similar to a quantum interactive proof system, with the exception that in the case of a positive

instance of the problem, the verifier should be able to generate the state on his systems with a polynomial-time quantum computer (so that he has not generated any state which he could have not have generated already on his own).[1] We prove this result by a reduction from the QSZK-complete promise problem QUANTUM-STATE-DISTINGUISHABILITY to QSEP-CIRCUIT. This reduction is somewhat similar to a previous reduction of Rosgen and Watrous [RW05], but the setting here is different and thus requires a different analysis (though interestingly, the analysis is reminiscent of arguments used in quantum information theory [ADHW09, Dev05]).

- QSEP-CIRCUIT is NP-hard. This result follows from the fact that the matrix version of the quantum separability problem is NP-hard, though we require some results of Knill which show that one can encode a quantum state efficiently by a unitary circuit if given a matrix description of the state [Kni95].

- We consider a variation of QSEP-CIRCUIT, called QSEP-CHANNEL, in which some of the inputs to the circuit are arbitrary and some are fixed in the state $|0\rangle$. The output qubits are again divided into three sets: the reference system qubits which are traced out, Alice's qubits, and Bob's qubits. We can think of this circuit as a quantum channel $\mathcal{N}_{S\to AB}$ with an input system S and two output systems A and B. The task is then to decide if there exists an input to the circuit such that the output state shared between Alice and Bob is separable, and we have the promise that either

$$\min_{\rho_S,\ \sigma_{AB}\in\mathcal{S}} \left\|\mathcal{N}_{S\to AB}\left(\rho_S\right)-\sigma_{AB}\right\|_1 \leq \varepsilon_1,$$

or

$$\min_{\rho_S,\ \sigma_{AB}\in\mathcal{S}} \left\|\mathcal{N}_{S\to AB}\left(\rho_S\right)-\sigma_{AB}\right\|_{1-\text{LOCC}} \geq \varepsilon_2,$$

where again the difference $\varepsilon_2-\varepsilon_1$ is larger than a sufficiently large inverse polynomial in the circuit size. We show that this promise problem is QIP-complete. The reasoning here is similar to the reasoning for our earlier results for QSEP-CIRCUIT, and we again exploit the results of Rosgen and Watrous [RW05] (in particular, the fact that QUANTUM-CIRCUIT-DISTINGUISHABILITY is QIP-complete). We will also refer to this promise problem throughout as the channel quantum separability problem.

- MULTI-QSEP-CIRCUIT, a multipartite generalization of QSEP-CIRCUIT, is also decidable by a two-message quantum interactive proof system for a wide range of parameters. The analysis of the proof system exploits a recent quantum de Finetti theorem of Brandão and Harrow [BH12] along with the analysis used in the proof system for QSEP-CIRCUIT. MULTI-QSEP-CIRCUIT is also NP- and QSZK-hard because QSEP-CIRCUIT trivially reduces to it, as QSEP-CIRCUIT is merely a special case of MULTI-QSEP-CIRCUIT.

The paper is structured as follows. In the next section, we introduce preliminary concepts such as quantum states and channels, distance measures between quantum states, k-extendibility, three- and two-message quantum interactive proof systems and complete promise problems for their corresponding complexity classes. Section 4 contains our first result, in which we demonstrate that

[1] This is the definition of QSZK in the case when the verifier behaves honestly. Watrous later showed that the expressive power of this class is the same as the case in which the verifier does not behave honestly [Wat09b].

a two-message quantum interactive proof system decides QSEP-CIRCUIT. We show in Section 5 that QSEP-CIRCUIT is hard for the class QSZK, and in Section 6, we show that it is NP-hard as well. In Section 7, we extend the above results to show that the channel variation of the quantum separability problem is QIP-complete. Section 8 gives a two-message quantum interactive proof system that decides MULTI-QSEP-CIRCUIT. Finally, we conclude in Section 9 with a summary and some open directions.

3 Preliminaries

In this preliminary section, we review background concepts from quantum information theory [NC00] and quantum computational complexity theory [Wat09a].

3.1 Quantum states and channels

A quantum state is a positive semidefinite, unit-trace operator (referred to as the density operator) acting on some Hilbert space $\mathcal{H}$. Let $\mathcal{D}(\mathcal{H})$ denote the set of density operators acting on a Hilbert space $\mathcal{H}$. An *extension* of a quantum state $\rho \in \mathcal{D}(\mathcal{H}_A)$ is some state $\omega \in \mathcal{D}(\mathcal{H}_A \otimes \mathcal{H}_B)$ (on a larger Hilbert space) such that $\rho = \mathrm{Tr}_{\mathcal{H}_B}\{\omega\}$. A quantum state is pure if its density operator is unit rank, in which case it has an equivalent representation as a unit vector $|\psi\rangle \in \mathcal{H}$. A *purification* of a density operator $\rho \in \mathcal{D}(\mathcal{H})$ is a pure extension of ρ. Throughout this work, we restrict ourselves to finite-dimensional Hilbert spaces, so that a d-dimensional Hilbert space is isomorphic to $\mathbb{C}^d$. A quantum channel is a linear, completely positive, trace-preserving (CPTP) map $\mathcal{N} : \mathcal{D}(\mathcal{H}_{\text{in}}) \to \mathcal{D}(\mathcal{H}_{\text{out}})$. The Stinespring representation theorem states that every CPTP map can be realized by tensoring its input with an ancillary environment system in some fiducial state $|0\rangle_E \in \mathcal{H}_E$ where $\dim(\mathcal{H}_E) \leq \dim(\mathcal{H}_{\text{in}})\dim(\mathcal{H}_{\text{out}})$, performing some unitary operation on the joint Hilbert space $\mathcal{H}_{\text{in}} \otimes \mathcal{H}_E$, factoring the unitary's output Hilbert space as $\mathcal{H}_{\text{out}} \otimes \mathcal{H}_{E'}$, and finally tracing over the Hilbert space $\mathcal{H}_{E'}$ [Sti55]. That is, for every CPTP map $\mathcal{N}$, there exists some unitary U such that the following relation holds for all $\rho \in \mathcal{D}(\mathcal{H}_{\text{in}})$:

$$\mathcal{N}(\rho) = \mathrm{Tr}_{E'}\left\{U\left(\rho \otimes |0\rangle\langle 0|_E\right)U^\dagger\right\}.$$

This theorem is the essential reason for the equivalence in computational power between the unitary and mixed-state circuit models of quantum computation [AKN98].

3.2 Distance measures

The trace norm of an operator A is $\|A\|_1 \equiv \mathrm{Tr}\{\sqrt{A^\dagger A}\}$. The metric on quantum states induced by the trace norm is called the trace distance, which has an operational interpretation as the bias when using an optimal measurement to distinguish two quantum states ρ and σ [NC00]. That is, when ρ and σ are chosen uniformly at random, the probability p_{succ} of successfully discriminating them with an optimal measurement is as follows:

$$p_{\text{succ}} = \frac{1}{2}\left(1 + \frac{1}{2}\|\rho - \sigma\|_1\right).$$

There is also a variational characterization of the trace distance as

$$\|\rho - \sigma\|_1 = 2 \max_{0 \leq \Lambda \leq I} \mathrm{Tr}\{\Lambda(\rho - \sigma)\},$$

which leads to the following useful inequality that holds for all Γ such that $0 \leq \Gamma \leq I$:

$$\mathrm{Tr}\{\Gamma\rho\} \geq \mathrm{Tr}\{\Gamma\sigma\} - \|\rho - \sigma\|_1 . \tag{5}$$

Suppose now that ρ_{AB} and σ_{AB} are in $\mathcal{D}(\mathcal{H}_A \otimes \mathcal{H}_B)$. Then the one-way local operations and classical communication (1-LOCC) distance between these two states, induced by a 1-LOCC norm [MWW09], is given by

$$\|\rho_{AB} - \sigma_{AB}\|_{1-\mathrm{LOCC}} \equiv \max_{\Lambda_{B\to X}} \|(I_A \otimes \Lambda_{B\to X})(\rho_{AB} - \sigma_{AB})\|_1 ,$$

where the maximization on the RHS is over all quantum-to-classical channels

$$\Lambda_{B\to X}(\omega) \equiv \sum_{x\in\mathcal{X}} \mathrm{Tr}\{\Lambda_x \omega\} |x\rangle\langle x|$$

with $\Lambda_x \geq 0$ for all $x \in \mathcal{X}$, $\sum_{x\in\mathcal{X}} \Lambda_x = I$, and $\{|x\rangle\}$ some orthonormal basis. (Note that we could also define the 1-LOCC distance with respect to measurement maps on Alice's system, which gives a value generally different from the above.) The 1-LOCC distance is equal to the bias in success probability when using an optimal one-way LOCC protocol to distinguish the states:

$$p_{\text{succ,1-LOCC}} = \frac{1}{2}\left(1 + \frac{1}{2}\|\rho_{AB} - \sigma_{AB}\|_{1-\mathrm{LOCC}}\right).$$

This distance is the natural distance measure in the setting of Bell experiments [Bel64] or quantum teleportation, for example. Also, it follows that

$$\|\rho - \sigma\|_{1-\mathrm{LOCC}} \leq \|\rho - \sigma\|_1 , \tag{6}$$

because a general protocol to distinguish ρ from σ never performs worse than one restricted to one-way LOCC operations.

The quantum fidelity $F(\rho, \sigma)$ between two quantum states ρ and σ is another measure of distinguishability, defined as follows:

$$F(\rho, \sigma) \equiv \left\|\sqrt{\rho}\sqrt{\sigma}\right\|_1^2 . \tag{7}$$

Uhlmann characterized the fidelity as the optimal squared overlap between any two purifications of ρ and σ [Uhl76]:

$$F(\rho, \sigma) = \max_{|\phi_\rho\rangle, |\phi_\sigma\rangle} |\langle\phi_\rho|\phi_\sigma\rangle|^2 .$$

Thus, the fidelity has an operational interpretation as the optimal probability with which a purification of ρ would pass a test for being a purification of σ. Since all purifications are related by a unitary operation on the purifying system, this characterization is equivalent to the following one:

$$F(\rho, \sigma) = \max_U |\langle\phi_\rho| (U \otimes I_{\mathcal{H}}) |\phi_\sigma\rangle|^2 , \tag{8}$$

where $|\phi_\rho\rangle$ and $|\phi_\sigma\rangle$ are now two fixed purifications of ρ and σ, respectively, and the optimization is over all unitaries acting on the purifying system (the fact that (7) is equal to (8) is often referred to as Uhlmann's theorem). The fidelity and trace distance are related by the Fuchs-van-de-Graaf inequalities [FvdG99]:

$$1 - \sqrt{F(\rho, \sigma)} \leq \frac{1}{2}\|\rho - \sigma\|_1 \leq \sqrt{1 - F(\rho, \sigma)}. \tag{9}$$

3.3 Separability, k-extendibility, and the maximum k-extendible fidelity

A bipartite state $\rho_{AB} \in \mathcal{D}(\mathcal{H}_A \otimes \mathcal{H}_B)$ is k-extendible [DPS02, DPS04] if there exists a state $\omega_{AB_1 \cdots B_k} \in \mathcal{D}(\mathcal{H}_A \otimes \mathcal{H}_{B_1} \otimes \cdots \otimes \mathcal{H}_{B_k})$ such that

1. Each Hilbert space $\mathcal{H}_{B_i}$ is isomorphic to $\mathcal{H}_B$ for all $i \in \{1, \ldots, k\}$.

2. The state $\omega_{AB_1 \cdots B_k}$ is invariant under permutations of the systems B_1 through B_k. That is,

$$\forall \pi \in S_k : \omega_{AB_1 \cdots B_k} = \left(I_A \otimes W^{\pi}_{B_1 \cdots B_k}\right) \omega_{AB_1 \cdots B_k} \left(I_A \otimes W^{\pi}_{B_1 \cdots B_k}\right)^{\dagger}, \tag{10}$$

   where S_k is the symmetric group on k elements and $W^{\pi}_{B_1 \cdots B_k}$ is a unitary operation that implements the permutation π of the B systems.

3. The state $\omega_{AB_1 \cdots B_k}$ is an extension of ρ_{AB}:

$$\rho_{AB} = \mathrm{Tr}_{B_2 \cdots B_k} \left\{\omega_{AB_1 \cdots B_k}\right\}.$$

Let $\mathcal{E}_k$ denote the set of k-extendible states. Every separable state is k-extendible for all $k \geq 2$. Since every separable state has a decomposition of the form in (2), an obvious choice for a k-extension is

$$\sum_{x \in \mathcal{X}} p_X(x) \; |\psi_x\rangle \langle\psi_x|_A \otimes |\phi_x\rangle \langle\phi_x|_{B_1} \otimes \cdots \otimes |\phi_x\rangle \langle\phi_x|_{B_k}. \tag{11}$$

On the other hand, if a state is not separable, there always exists some k for which the state is not k-extendible, and furthermore, for every $l > k$, the state is also not l-extendible [DPS02, DPS04]. In this sense, the set $\mathcal{E}_k$ forms an approximation to the set $\mathcal{S}$ of separable states, and the approximation becomes exact in the limit as $k \to \infty$.

The *maximum k-extendible fidelity* of a state ρ_{AB} is defined as

$$\max_{\sigma_{AB} \in \mathcal{E}_k} F(\rho_{AB}, \sigma_{AB}).$$

In this paper, we give the maximum k-extendible fidelity an operational interpretation as an upper bound on the maximum probability with which a prover can convince a verifier to accept in our $\mathsf{QIP(2)}$ protocol that tests for k-extendibility. Clearly, the above quantity converges to the *maximum separable fidelity* (defined in [Wat04]) in the limit as $k \to \infty$:

$$\max_{\sigma_{AB} \in \mathcal{S}} F(\rho_{AB}, \sigma_{AB}) = \lim_{k \to \infty} \max_{\sigma_{AB} \in \mathcal{E}_k} F(\rho_{AB}, \sigma_{AB}).$$

Finally, we state a lemma which extends Theorem 3 of [BCY11b]. This lemma establishes a notion of approximate k-extendibility which is essential for our work here. The proof of this lemma is a straightforward modification of the proof of Theorem 3 of [BCY11b] and we provide it in the appendix.

Lemma 1 *Let ρ_{AB} be ε-away in one-way LOCC distance from the set of separable states, for some $\varepsilon > 0$:*

$$\min_{\sigma_{AB} \in \mathcal{S}} \left\|\rho_{AB} - \sigma_{AB}\right\|_{1\text{-}LOCC} \geq \varepsilon.$$

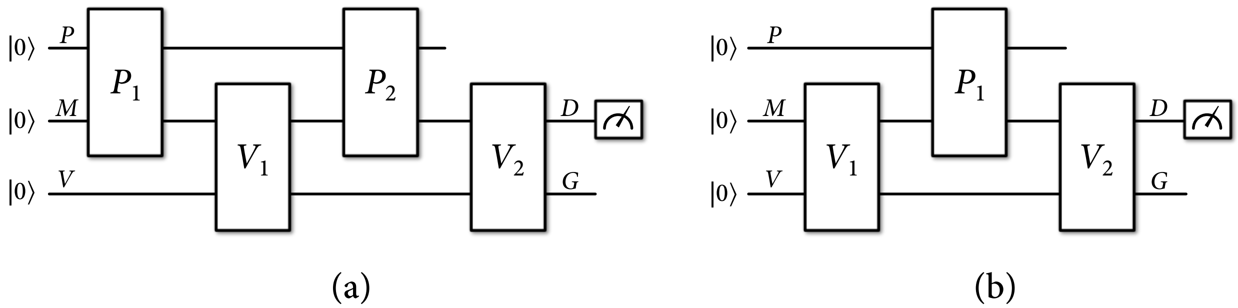

Figure 2: Quantum interactive proof systems with (a) three messages and (b) two messages exchanged between the verifier and the prover.

Then the state ρ_{AB} is δ-away in trace distance from the set of k-extendible states:

$$\min_{\sigma_{AB}\in\mathcal{E}_k} \|\rho_{AB} - \sigma_{AB}\|_1 \geq \delta,$$

for $\delta < \varepsilon$ and where

$$k = \left\lceil \frac{16 \ln 2 \ \log|A|}{(\varepsilon - \delta)^2} \right\rceil.$$

3.4 Quantum interactive proof systems

Quantum interactive proof systems were first defined and analyzed in [Wat02, KW00]. Their study is an important component of quantum computational complexity theory, and they are essential in our work here. A quantum interactive proof system involves an exchange of quantum messages between a polynomial-time quantum verifier and a computationally unbounded quantum prover. Since any quantum interactive proof system can be parallelized such that at most three messages are exchanged between the verifier and prover ($\mathsf{QIP} = \mathsf{QIP}(3)$ up to a wide range of parameters) [KW00], we focus our attention on quantum interactive proof systems with three or fewer messages, starting by defining $\mathsf{QIP}(3)$.

3.4.1 Three-message quantum interactive proof systems

Given an input string x of length n, a quantum verifier consists of two unitary quantum circuits V_1 and V_2 computed from x in polynomial time. These circuits are generated from some finite gate set that is universal for quantum computation [NC00, Wat09a]. The qubits of the verifier are divided into two sets: private qubits and message qubits. The verifier exchanges the message qubits with the prover, while keeping the private qubits in his own laboratory. A quantum prover is defined similarly to the verifier—he has private qubits and message qubits that he exchanges with the verifier. However, he is allowed to perform two arbitrary, unconstrained unitary quantum operations P_1 and P_2 on the qubits in his laboratory.

Figure 2(a) depicts a quantum interactive proof system with three messages. P is the register containing the prover's private qubits, M contains the message qubits, and V contains the verifier's private qubits. For simplicity, we take the convention that these registers may change size throughout the execution of the protocol, as long as registers M and V are polynomial in the length of

the input x. The quantum interactive proof system begins with the prover initializing the qubits in his registers P and M to the all-zero state, and the verifier does the same to his qubits in the register V. The prover then acts with the unitary P_1 on registers P and M and sends the message qubits to the verifier. They proceed along the lines indicated in Figure 2(a), until the verifier acts with a final unitary V_2. This final unitary has two output systems: a decision qubit in register D and other "garbage qubits" in register G. The verifier then measures the decision qubit in the computational basis to decide whether to accept or reject. The maximum probability with which the prover can make the verifier accept is equal to

$$\max_{|\psi\rangle_{PM},\ P_2} \left\| \langle 1|_D \, V_2 P_2 V_1 \, |\psi\rangle_{PM} \, |0\rangle_V \right\|_2^2 , \tag{12}$$

where the maximization is over all pure states $|\psi\rangle_{PM} = P_1 \, |0\rangle_{PM}$ that the prover can prepare at the beginning of the protocol and all unitaries P_2 that he can perform.

By defining $\mathcal{N}_1$ as the quantum channel induced by applying the verifier's first unitary and tracing over the verifier's message register M:

$$\mathcal{N}_1(\rho) \equiv \mathrm{Tr}_M \left\{ V_1 \left(\rho \otimes |0\rangle \langle 0|_V \right) V_1^\dagger \right\} ,$$

defining $\mathcal{N}_2$ as the quantum channel induced by applying the inverse of the verifier's final unitary and tracing over the verifier's message register M:

$$\mathcal{N}_2(\sigma) \equiv \mathrm{Tr}_M \left\{ V_2^\dagger \left(\sigma_G \otimes |1\rangle \langle 1|_D \right) V_2 \right\} , \tag{13}$$

and applying Uhlmann's theorem in (8) to the maximum acceptance probability in (12), one can rewrite (12) as the maximum output fidelity of the channels $\mathcal{N}_1$ and $\mathcal{N}_2$ over all possible inputs to them [KW00, RW05, Ros09]:

$$\max_{|\psi\rangle_{PM},\ P_2} \left\| \langle 1|_D \, V_2 P_2 V_1 \, |\psi\rangle_{PM} \, |0\rangle_V \right\|_2^2 = \max_{\rho,\ \sigma} F\left(\mathcal{N}_1(\rho), \mathcal{N}_2(\sigma) \right) . \tag{14}$$

(See Chapter 4 of [Ros09] for details of this calculation).

We can now define the complexity class QIP(3):

Definition 2 (QIP(3)) *Let $A = (A_{yes}, A_{no})$ be a promise problem, and let $c, s : \mathbb{N} \to [0,1]$ be polynomial-time computable functions such that the gap $c - s$ is at least an inverse polynomial in the input size. Then $A \in$ QIP$(3, c, s)$ if there exists a three-message quantum interactive proof system with the following properties:*

1. *Completeness: For all input strings $x \in A_{yes}$, there exists a prover that causes the verifier to accept with probability at least $c(|x|)$.*
2. *Soundness: For all input strings $x \in A_{no}$, every prover causes the verifier to accept with probability at most $s(|x|)$.*

Note that one can amplify the gap between c and s to a constant by employing parallel repetition [KW00].

Given the fact that we can rewrite the maximum acceptance probability of any QIP(3) protocol as in (14), it is straightforward to formulate a complete promise problem for QIP(3) called CLOSE-IMAGES, which essentially just amounts to rewriting the above definition [RW05, Ros09]:

Problem 3 (CLOSE-IMAGES) *Fix two constants $c, s \in [0,1]$ such that $c > s$. Given are two mixed-state quantum circuits Q_0 and Q_1, each accepting n-qubit inputs and having m-qubit outputs. Decide whether*

1. *Yes: There exist n-qubit states ρ_0 and ρ_1 such that $\max_{\rho_0,\rho_1} F\left(Q_0\left(\rho_0\right), Q_1\left(\rho_1\right)\right) \geq c$.*
2. *No: For all n-qubit states ρ_0 and ρ_1, it holds that $\max_{\rho_0,\rho_1} F\left(Q_0\left(\rho_0\right), Q_1\left(\rho_1\right)\right) \leq s$.*

The following promise problem is also complete for QIP(3), but proving so requires more than a trivial rewriting of the definition of QIP(3) [RW05, Ros09]:

Problem 4 (QUANTUM-CIRCUIT-DISTINGUISHABILITY) *Fix a constant $\varepsilon \in [0,1)$. Given are two mixed-state quantum circuits Q_0 and Q_1, each with n-qubit inputs and m-qubit outputs. Decide whether*

1. *Yes: There is a quantum input for which the circuits are distinguishable:*

$$\max_{\rho\in\mathcal{D}(\mathcal{H}_R\otimes\mathcal{H}_{in})} \left\|\left(I_R\otimes Q_0\right)(\rho) - \left(I_R\otimes Q_1\right)(\rho)\right\|_1 \geq 2-\varepsilon.$$

2. *No: No quantum input can distinguish the circuits:*

$$\max_{\rho\in\mathcal{D}(\mathcal{H}_R\otimes\mathcal{H}_{in})} \left\|\left(I_R\otimes Q_0\right)(\rho) - \left(I_R\otimes Q_1\right)(\rho)\right\|_1 \leq \varepsilon.$$

In what follows, we abbreviate QUANTUM-CIRCUIT-DISTINGUISHABILITY as QCD.

3.4.2 Two-message quantum interactive proof systems

Now that we have defined three-message quantum interactive proof systems, it is straightforward to define two-message quantum interactive proof systems. The description is essentially identical to that given in the first two paragraphs of the previous section, with the exception that the verifier acts both first and last with unitaries V_1 and V_2, while the prover acts between these two operations with a unitary P_1 (see Figure 2(a)). Thus, the maximum probability with which the prover can make the verifier accept is equal to

$$\max_{P_2} \left\|\langle 1|_D\, V_2 P_2 \,|0\rangle_P \,|\phi\rangle_{MV}\right\|_2^2,$$

where $|\phi\rangle_{MV} = V_1\,|0\rangle_{MV}$. By following essentially the same reasoning as before, we can rewrite this maximum acceptance probability in terms of the quantum fidelity. First, we define the mixed state

$$\omega_V = \mathrm{Tr}_M\left\{V_1\ \ |0\rangle\langle 0|_{MV}\ \ V_1^\dagger\right\}.$$

We also define the quantum channel $\mathcal{N}$ from V_2 as in (13). By again applying Uhlmann's theorem, it is straightforward to prove that the maximum acceptance probability is equal to the fidelity between the state ω and the channel $\mathcal{N}$ when maximizing over all inputs to the channel:

$$\max_{P_2} \left\|\langle 1|_D\, V_2 P_2\,|0\rangle_P\,|\phi\rangle_{MV}\right\|_2^2 = \max_{\sigma} F\left(\omega, \mathcal{N}\left(\sigma\right)\right). \tag{15}$$

The definition of the complexity class QIP(2) is identical to that given in Definition 2 (but substituting QIP(2) for QIP(3)). One can similarly amplify the gap between c and s to a constant by employing parallel repetition [JUW09]. Also, the following promise problem, called CLOSE-IMAGE, is a complete promise problem for QIP(2):

Problem 5 (CLOSE-IMAGE) *Fix two constants $c, s \in [0,1]$ such that $c > s$. Given is a mixed-state quantum circuit to generate the m-qubit state ρ_0 and a mixed-state quantum circuit Q_1, with an n-qubit input state and an m-qubit output state. Decide whether*

1. *Yes: There exists an n-qubit state ρ_1 such that $\max_{\rho_1} F(\rho_0, Q_1(\rho_1)) \geq c$.*
2. *No: For all n-qubit states ρ_1, it holds that $\max_{\rho_1} F(\rho_0, Q_1(\rho_1)) \leq s$.*

The fact that CLOSE-IMAGE is complete for QIP(2) follows essentially the same reasoning as in [Ros09] (it amounts to a rewriting of the definition of QIP(2)).

3.4.3 Quantum statistical zero-knowledge proof systems

Another kind of quantum interactive proof system that is relevant for our work here is a quantum statistical zero-knowledge (QSZK) proof system [Wat02, Wat09b]. The definition of the "honest-verifier" version of this complexity class is essentially the same as that for QIP (with an arbitrary number of messages exchanged), but the difference is that in the case of a positive problem instance, the states of the verifier before and after every interaction with the prover should be such that he could have actually generated them himself. In this sense, he does not gain any "knowledge" by interacting with the prover (other than being convinced to accept). Watrous has shown that any honest-verifier QSZK proof system has an equivalent proof system in which the verifier is not required to behave honestly [Wat09b].

The following promise problem is complete for QSZK [Wat02]:

Problem 6 (QUANTUM-STATE-DISTINGUISHABILITY) *Fix a constant $\varepsilon \in [0,1)$. Given is a mixed-state quantum circuit to generate the n-qubit states ρ_0 and ρ_1. Decide whether*

1. *Yes: $\|\rho_0 - \rho_1\|_1 \geq 2 - \varepsilon$.*
2. *No: $\|\rho_0 - \rho_1\|_1 \leq \varepsilon$.*

In what follows, we abbreviate QUANTUM-STATE-DISTINGUISHABILITY as QSD.

4 A two-message quantum interactive proof system decides the quantum separability problem

We are now ready to provide a formal statement of the promise problem QSEP-CIRCUIT, and we follow with a proof that it is decidable by a two-message quantum interactive proof system.

Problem 7 (QSEP-CIRCUIT(δ_c, δ_s)) *Given is a mixed-state quantum circuit to generate the n-qubit state ρ_{AB}, along with a labeling of the qubits in the reference system R and the output qubits for A and B. Decide whether*

1. *Yes: There is a separable state $\sigma_{AB} \in \mathcal{S}$ that is δ_c-close to ρ_{AB} in trace distance:*

$$\min_{\sigma_{AB} \in \mathcal{S}} \|\rho_{AB} - \sigma_{AB}\|_1 \leq \delta_c.$$

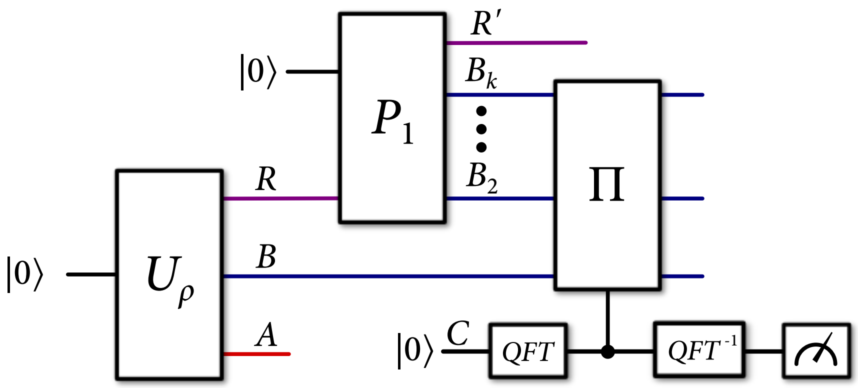

Figure 3: A two-message quantum interactive proof system for QSEP-CIRCUIT. It begins with the verifier executing the circuit U_ρ that generates the state ρ_{AB}. He sends the reference system to the prover. In the case that ρ_{AB} is separable, the prover should be able to act with a unitary on the reference system and some ancillas in order to generate a purification of a k-extension of ρ_{AB}. The prover sends all of the extension systems back to the verifier, who then performs phase estimation over the symmetric group (a quantum Fourier transform followed by a controlled permutation) in order to test if the state sent by the prover is really a k-extension.

2. *No: Every separable state is at least δ_s-far from ρ_{AB} in 1-LOCC distance:*

$$\min_{\sigma_{AB}\in\mathcal{S}} \|\rho_{AB} - \sigma_{AB}\|_{1-LOCC} \geq \delta_s.$$

Theorem 8 *QSEP-CIRCUIT$(\delta_c, \delta_s) \in$ QIP(2) if there are polynomial-time computable functions $\delta_c, \delta_s : \mathbb{N} \to [0,1]$, such that the difference $\delta_s^2/8 - 2\sqrt{\delta_c}$ is larger than an inverse polynomial in the circuit size.*

Proof. Figure 3 depicts a two-message quantum interactive proof system for QSEP-CIRCUIT. The protocol begins with the verifier preparing the state $|\psi_\rho\rangle_{RAB}$, a particular purification of ρ_{AB}, by running the quantum circuit U_ρ specified by the input string (the problem instance). The verifier transmits the reference system to the prover, who then acts on R and some ancillary qubits with a unitary P_1 that has output systems R', B_2, …, B_k. The prover transmits systems B_2, …, B_k to the verifier. The verifier then performs phase estimation over the symmetric group [Kit95, BBD+97] (also known as the "permutation test" [KNY08]) on the registers B, B_2, …, B_k, using the qubits in system C as the control. The verifier performs a computational basis measurement on all of the qubits in the control register C and accepts if and only if the measurement outcomes are all zeros.

This protocol is just an implementation of a k-extendibility test on a quantum computer. We can build intuition for why it works on YES instances by examining the exact case, when ρ_{AB} is actually a separable state. In this case, we know that ρ_{AB} has a decomposition of the form given in (2), and as such, it has an extension of the form in (11) for all $k > 1$. Thus, the following state is a purification of ρ_{AB}:

$$|\phi_{k,\rho}\rangle_{R'ABB_2\cdots B_k} \equiv \sum_{x\in\mathcal{X}} \sqrt{p_X(x)}\ |x\rangle_{R'} \otimes |\psi_x\rangle_A \otimes |\phi_x\rangle_B \otimes |\phi_x\rangle_{B_2} \otimes \cdots \otimes |\phi_x\rangle_{B_k},$$

where $\{|x\rangle_{R'}\}$ is some orthonormal basis for the reference system. Since all purifications are related by unitaries on the reference system, the prover can append ancilla qubits to the R system received from the verifier and perform a unitary P_1 that takes $|\psi_\rho\rangle_{RAB}|0\rangle$ to $|\phi_{k,\rho}\rangle_{R'ABB_2\cdots B_k}$. The prover then sends the systems $B_2, \ldots, B_k$ to the verifier. The verifier performs a permutation test on the systems $B, B_2, \ldots, B_k$. Since the state $|\phi_{k,\rho}\rangle_{R'ABB_2\cdots B_k}$ is invariant under permutations of the systems $B, B_2, \ldots, B_k$, the qubits in the control register C do not acquire a phase. Thus, after the final quantum Fourier transform is applied, the qubits in the control register C are in the all-zero state with certainty.

The analysis for a YES instance follows the above intuition closely. In this case, there is some state $\sigma_{AB} \in \mathcal{S}$ that is δ_c-close in trace distance to ρ_{AB}. By Uhlmann's theorem in (8) and the Fuchs-van-de-Graaf inequalities in (9), there is a purification $|\psi_\sigma\rangle_{RAB}$ of σ_{AB} such that

$$\left\| |\psi_\rho\rangle\langle\psi_\rho|_{RAB} - |\psi_\sigma\rangle\langle\psi_\sigma|_{RAB} \right\|_1 \leq 2\sqrt{\delta_c}. \tag{16}$$

So the prover can just operate as above, but choosing his unitary P_1 to correspond to the state $|\psi_\sigma\rangle_{RAB}$ instead. Writing as U the unitary corresponding to P_1 followed by the permutation test, we obtain the following lower bound on the probability with which the verifier accepts:

$$\begin{aligned}
&\operatorname{Tr}\left\{|0\rangle\langle 0|_C\, U\left(|\psi_\rho\rangle\langle\psi_\rho|_{RAB}\right)U^\dagger\right\} \\
&= \operatorname{Tr}\left\{U^\dagger|0\rangle\langle 0|_C\, U\left(|\psi_\rho\rangle\langle\psi_\rho|_{RAB}\right)\right\} \\
&\geq \operatorname{Tr}\left\{U^\dagger|0\rangle\langle 0|_C\, U\left(|\psi_\sigma\rangle\langle\psi_\sigma|_{RAB}\right)\right\} - \left\| |\psi_\rho\rangle\langle\psi_\rho|_{RAB} - |\psi_\sigma\rangle\langle\psi_\sigma|_{RAB} \right\|_1 \\
&\geq 1 - 2\sqrt{\delta_c},
\end{aligned} \tag{17}$$

where the first inequality follows from (5), and the second inequality follows by applying (16) and because the protocol accepts with probability one for a separable state.

The analysis for a NO instance has two components:

1. demonstrating that the maximum k-extendible fidelity is an upper bound on the maximum acceptance probability
2. using Lemma 1 regarding approximate k-extendibility and the first item above to specify how large k should be in order to obtain a good upper bound on the maximum acceptance probability.[2]

We now discuss the first item above. Recall from (15) that the maximum acceptance probability of any QIP(2) system is equal to the maximum fidelity between the state generated by the verifier's first circuit and the channel generated by the inverse of the verifier's second circuit. For the protocol in Figure 3, the state generated by the verifier's first circuit is as follows:

$$\rho_{AB} \otimes |\text{perm}\rangle\langle\text{perm}|_C\,,$$

where $|\text{perm}\rangle_C$ is a superposition over all possible permutations of k elements resulting from an application of the quantum Fourier transform [NC00] to the state $|0\rangle_C$:

$$|\text{perm}\rangle_C \equiv \frac{1}{\sqrt{k!}}\sum_{\pi\in S_k}|\pi\rangle_C\,, \tag{18}$$

[2]For a YES instance, the value of k does not matter because the lower bound on the maximum acceptance probability is always as given above.

so that the C register requires $\lceil \log_2(k!) \rceil$ qubits. (Note that Figure 3 depicts the verifier generating $|\text{perm}\rangle_C$ later in the protocol, but we could just as easily reorder things so that he generates this state in the first step.) The channel generated by the inverse of the verifier's circuit conditional on accepting is

$$\mathcal{M}_{ABB_2\cdots B_k \to ABC}(\sigma_{ABB_2\cdots B_k}) \equiv \text{Tr}_{B_2\cdots B_k}\left\{ (U_\Pi)_{BB_2\cdots B_k C} (\sigma_{ABB_2\cdots B_k} \otimes |\text{perm}\rangle\langle\text{perm}|_C) \left(U_\Pi^\dagger\right)_{BB_2\cdots B_k C} \right\},$$

where $(U_\Pi)_{BB_2\cdots B_k C}$ is a controlled-permutation operation:

$$(U_\Pi)_{BB_2\cdots B_k C} \equiv \sum_{\pi \in S_k} W^\pi_{BB_2\cdots B_k} \otimes |\pi\rangle\langle\pi|_C, \tag{19}$$

and $W^\pi_{BB_2\cdots B_k}$ is a unitary operation corresponding to permutation π (mentioned before in (10)). So the maximum acceptance probability is equal to

$$\max_{\sigma_{ABB_2\cdots B_k}} F(\rho_{AB} \otimes |\text{perm}\rangle\langle\text{perm}|_C, \mathcal{M}_{ABB_2\cdots B_k \to ABC}(\sigma_{ABB_2\cdots B_k})).$$

Since the fidelity can only increase under the discarding of the control register C,[3] the maximum acceptance probability is upper bounded by the following quantity:

$$\max_{\sigma_{ABB_2\cdots B_k}} F(\rho_{AB}, \mathcal{M}_{ABB_2\cdots B_k \to AB}(\sigma_{ABB_2\cdots B_k})), \tag{20}$$

where

$$\begin{aligned}\mathcal{M}_{ABB_2\cdots B_k \to AB}(\sigma_{ABB_2\cdots B_k}) &= \text{Tr}_C\{\mathcal{M}_{ABB_2\cdots B_k \to ABC}(\sigma_{ABB_2\cdots B_k})\} \\ &= \frac{1}{k!}\sum_{\pi \in S_k} \text{Tr}_{B_2\cdots B_k}\left\{ \left(I_A \otimes W^\pi_{BB_2\cdots B_k}\right) \sigma_{ABB_2\cdots B_k} \left(I_A \otimes W^\pi_{BB_2\cdots B_k}\right)^\dagger \right\},\end{aligned}$$

which is just the channel that applies a random permutation of the B systems and discards the last $k-1$ systems $B_2, \ldots, B_k$. Clearly, since the channel $\mathcal{M}_{ABB_2\cdots B_k \to AB}$ symmetrizes the state of the systems $BB_2\cdots B_k$, the maximum in (20) is achieved by a state $\sigma_{ABB_2\cdots B_k}$ for which systems $BB_2\cdots B_k$ are permutation symmetric. Thus, by recalling the definition of k-extendibility, we can rewrite (20) as the maximum k-extendible fidelity of ρ_{AB}:

$$\max_{\sigma_{ABB_2\cdots B_k}} F(\rho_{AB}, \mathcal{M}_{ABB_2\cdots B_k \to AB}(\sigma_{ABB_2\cdots B_k})) = \max_{\sigma_{AB} \in \mathcal{E}_k} F(\rho_{AB}, \sigma_{AB}). \tag{21}$$

This demonstrates that the maximum k-extendible fidelity is an upper bound on the maximum acceptance probability and completes our proof of the first item above.

The second part of the analysis of a NO instance involves determining how large k needs to be. Suppose that

$$\min_{\sigma_{AB} \in \mathcal{S}} \|\rho_{AB} - \sigma_{AB}\|_{\text{1-LOCC}} \geq \delta_s.$$

[3]We can interpret discarding the control register as actually giving it to the prover, so that the resulting fidelity corresponds to the maximum acceptance probability in a modified protocol in which the prover controls the inputs to C.

According to Lemma 1, if we take k to be larger than

$$\left\lceil \frac{16 \ln 2 \ \log |A|}{(\delta_s - \delta_s')^2} \right\rceil .$$

then we can guarantee that

$$\min_{\sigma_{AB} \in \mathcal{E}_k} \left\| \rho_{AB} - \sigma_{AB} \right\|_1 \geq \delta_s'.$$

for δ_s' strictly less than δ_s. We can enforce this latter condition by setting $\delta_s' = \delta_s/\sqrt{2}$. Observe that k is polynomial in n_A, where n_A is the number of qubits in Alice's system. Then using the following manipulation of the Fuchs-van-de-Graaf inequalities in (9):

$$F(\rho, \sigma) \leq 1 - \frac{1}{4} \left\| \rho - \sigma \right\|_1^2 ,$$

we have that

$$\max_{\sigma_{AB} \in \mathcal{E}_k} F(\rho_{AB}, \sigma_{AB}) \leq 1 - \frac{1}{4} \min_{\sigma_{AB} \in \mathcal{E}_k} \left\| \rho_{AB} - \sigma_{AB} \right\|_1^2 \tag{22}$$

$$\leq 1 - \frac{1}{4} \left(\delta_s' \right)^2 \tag{23}$$

$$= 1 - \delta_s^2 / 8 \tag{24}$$

In the above, we have separated the probability of accepting and the probability of rejecting by an inverse polynomial in n_A (namely, from the promise that the difference $\delta_s^2/8 - 2\sqrt{\delta_c}$ is at least an inverse polynomial in the circuit size), and it is known that an inverse polynomial gap is sufficient to place this protocol in QIP(2) (see Section 3.2 of Ref. [JUW09] for how to amplify an inverse polynomial gap). Thus, we have given a two-message quantum interactive proof system that decides the quantum separability problem. ■

If one considers the characterization of QIP(2) in terms of the complete promise problem CLOSE-IMAGE, one has to identify a state and a channel to compare by means of the fidelity maximized over all inputs to the channel. The natural state to consider for QSEP-CIRCUIT is ρ_{AB}, while the natural channel to test for k-extendibility is one that applies a random permutation to the systems B, B_2, …, B_k of a state $\sigma_{ABB_2 \cdots B_k}$ and traces out the systems B_2, …, B_k. The Stinespring dilation of the channel is the circuit for phase estimation over the symmetric group, so that in this sense, we can say that CLOSE-IMAGE finds the phase estimation circuit given in Figure 3.

Also, from the above proof, it is clear that if the promise regarding QSEP-CIRCUIT is given in terms of fidelities rather than the trace distance and the 1-LOCC distance, then the promise concerning the gap between δ_s and δ_c could be that their difference is larger than an inverse polynomial (rather than making a promise about the difference $\delta_s^2/8 - 2\sqrt{\delta_c}$).

5 QSZK-hardness of the quantum separability problem

Having placed an upper bound on the difficulty of solving QSEP-CIRCUIT, we now move on to lower bounds, beginning in this section with a proof that it is hard for QSZK. Our approach is to exhibit a Karp reduction from the QSZK-complete promise problem QSD to QSEP-CIRCUIT. The essential idea behind the reduction is similar to Rosgen and Watrous's reduction of CLOSE-IMAGES to QCD [RW05, Ros09].

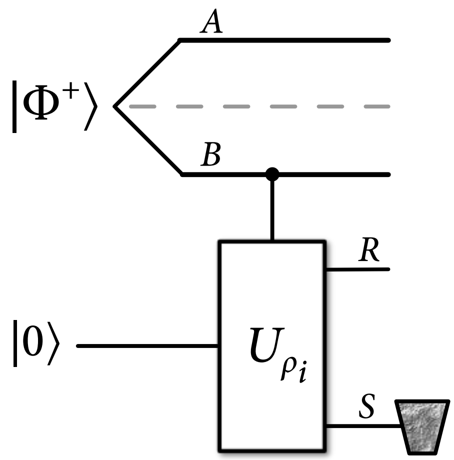

Figure 4: Given respective circuit descriptions U_{ρ_0} and U_{ρ_1} for generating the states ρ_0 and ρ_1 on the output system S, one can compute a description for the above circuit in polynomial time, and furthermore, the above circuit can be run efficiently on a quantum computer. This serves as a reduction from QUANTUM-STATE-DISTINGUISHABILITY to QSEP-CIRCUIT, i.e., where one should decide if the state on systems A and BR is separable with respect to this cut.

In order to demonstrate this reduction, we have to show that there is a polynomial-time algorithm that encodes YES instances of QSD into YES instances of QSEP-CIRCUIT and the same for the NO instances. Recall that for QSD, we are given a description of circuits U_{ρ_0} and U_{ρ_1} that generate mixed states ρ_0 and ρ_1. The output qubits of the circuit are divided into two sets: qubits in a reference system R that are traced over and qubits in a system S which contains ρ_i. For $i \in \{0,1\}$, let

$$|\psi_{\rho_i}\rangle_{RS} \equiv U_{\rho_i} |0\rangle ,$$

so that

$$\rho_i = \mathrm{Tr}_R \left\{ |\psi_{\rho_i}\rangle \langle\psi_{\rho_i}|_{RS} \right\} .$$

Figure 4 depicts a circuit that accomplishes the reduction. From the description of the circuits U_{ρ_0} and U_{ρ_1}, one can generate a description of the circuit in Figure 4 in polynomial time, and furthermore, the resulting circuit runs efficiently on a quantum computer [Ros09]. The circuit takes as input a Bell state

$$\left|\Phi^+\right\rangle_{AB} \equiv \frac{1}{\sqrt{2}} \left(|00\rangle_{AB} + |11\rangle_{AB} \right),$$

and performs the following controlled unitary from the qubit B to the ancilla qubits:

$$|0\rangle \langle 0|_B \otimes U_{\rho_0} + |1\rangle \langle 1|_B \otimes U_{\rho_1}.$$

The resulting state is as follows:

$$|\varphi\rangle_{ABRS} \equiv \frac{1}{\sqrt{2}} \left(|0\rangle_A |0\rangle_B |\psi_{\rho_0}\rangle_{RS} + |1\rangle_A |1\rangle_B |\psi_{\rho_1}\rangle_{RS} \right).$$

The output qubits are divided into three sets: environment qubits in the system S that are traced over, a single qubit in system A, and qubits in systems BR. Thus, the state resulting from applying the circuit in Figure 4 is as follows:

$$\omega_{A:BR} \equiv \mathrm{Tr}_S \left\{ |\varphi\rangle \langle\varphi|_{ABRS} \right\}. \tag{25}$$

The task is to decide whether the state on systems A and BR is separable across this cut, subject to the promise in Problem 7. Our claim is that YES instances of QSD map to YES instances of QSEP-CIRCUIT, with the same holding true for NO instances.

The intuition for why this reduction works is as follows. In the case of a YES instance of QSD, the states ρ_0 and ρ_1 are approximately orthogonal, so that tracing out the S system of the circuit in Figure 4 decoheres the Bell state, leaving a state on A and BR close to the following state:

$$\omega^{\mathrm{sep}}_{A:BR} \equiv \frac{1}{2} \left(|0\rangle \langle 0|_A \otimes |0\rangle \langle 0|_B \otimes (\psi_{\rho_0})_R + |1\rangle \langle 1|_A \otimes |1\rangle \langle 1|_B \otimes (\psi_{\rho_1})_R \right). \tag{26}$$

The above state is clearly separable with respect to the bipartite cut $A : BR$. In the case of a NO instance of QSD, the states ρ_0 and ρ_1 are approximately indistinguishable, and tracing over the S system of the circuit in Figure 4 does little to decohere the entanglement shared between A and BR. Thus, Bob can perform a local unitary operation on systems B and R to distill out a pure Bell state shared between A and B. After this, Alice and Bob can perform a Bell experiment on the distilled Bell state to determine if they indeed share a Bell state. Since these two operations can be performed with local operations and one message of classical communication, the resulting state is 1-LOCC distinguishable from the set of separable states.

We now give a formal proof to justify this reduction:

Theorem 9 *QSEP-CIRCUIT with constant promise gap is QSZK-hard.*

Proof. We first prove that the circuit in Figure 4 maps YES instances of QSD to YES instances of QSEP-CIRCUIT. So we begin by assuming that

$$\|\rho_0 - \rho_1\|_1 \geq 2 - \varepsilon, \tag{27}$$

and we will use this condition to show that the fidelity between $\omega^{\mathrm{sep}}_{A:BR}$ in (26) and the reduced state $\omega_{A:BR}$ in (25) is close to one. So, recall from Uhlmann's theorem in (8) that the fidelity between $\omega^{\mathrm{sep}}_{A:BR}$ and $\omega_{A:BR}$ is the maximum squared overlap between any purifications of these states. Thus, if we can show that the squared overlap between two *particular* purifications of $\omega^{\mathrm{sep}}_{A:BR}$ and $\omega_{A:BR}$ is large, then this implies a lower bound on the fidelity between these two states. Consider the following particular purification of $\omega^{\mathrm{sep}}_{A:BR}$:

$$|\omega^{\mathrm{sep}}_{ABB'RS}\rangle \equiv \frac{1}{\sqrt{2}} \left(|0\rangle_A |0\rangle_B |0\rangle_{B'} |\psi_{\rho_0}\rangle_{RS} + |1\rangle_A |1\rangle_B |1\rangle_{B'} |\psi_{\rho_1}\rangle_{RS} \right).$$

Recall that the trace distance bound in (27) implies the existence of a two-outcome projective measurement $\{\Pi_0, \Pi_1\}$ (known as a Helstrom measurement [Hel69, Hol72, Hel76]) that has the following success probability in discriminating ρ_0 from ρ_1 if they are chosen uniformly at random:

$$\begin{aligned} \frac{1}{2}\mathrm{Tr}\left\{\Pi_0 \rho_0\right\} + \frac{1}{2}\mathrm{Tr}\left\{\Pi_1 \rho_1\right\} &= \frac{1}{2}\left(1 + \frac{1}{2}\|\rho_0 - \rho_1\|_1\right) \\ &\geq 1 - \frac{\varepsilon}{2}. \end{aligned} \tag{28}$$

Performing the following "Helstrom isometry"

$$U^H_{S\to SB'} \equiv (\Pi_0)_S \otimes |0\rangle_{B'} + (\Pi_1)_S \otimes |1\rangle_{B'}$$

on the S system of $|\varphi\rangle_{ABRS}$ produces a particular purification of the state $\omega_{A:BR}$:

$$U^H_{S\to SB'} |\varphi\rangle_{ABRS} = \frac{1}{\sqrt{2}} \sum_{i,j\in\{0,1\}} |i\rangle_A |i\rangle_B \otimes |j\rangle_{B'} \otimes (\Pi_j)_S |\psi_{\rho_i}\rangle_{RS}.$$

The overlap between these purifications is

$$\begin{aligned}
&\langle \omega^{\text{sep}}_{ABB'RS} | U^H_{S\to SB'} |\varphi\rangle_{ABRS} \\
&= \frac{1}{2}\left(\sum_{k\in\{0,1\}} \langle k|_A \langle k|_B \langle k|_{B'} \langle \psi_{\rho_k}|_{RS}\right)\left(\sum_{i,j\in\{0,1\}} |i\rangle_A |i\rangle_B \otimes |j\rangle_{B'} \otimes (\Pi_j)_S |\psi_{\rho_i}\rangle_{RS}\right) \\
&= \frac{1}{2}\sum_{i,j,k\in\{0,1\}} \langle k|i\rangle_A \ \langle k|i\rangle_B \ \langle k|j\rangle_{B'} \ \langle \psi_{\rho_k}|_{RS}\, I_R \otimes (\Pi_j)_S |\psi_{\rho_i}\rangle_{RS} \\
&= \frac{1}{2}\sum_{i\in\{0,1\}} \langle \psi_{\rho_i}|_{RS}\, I_R \otimes (\Pi_i)_S |\psi_{\rho_i}\rangle_{RS} \\
&= \frac{1}{2}\mathrm{Tr}\{\Pi_0\rho_0\} + \frac{1}{2}\mathrm{Tr}\{\Pi_1\rho_1\} \\
&\geq 1 - \frac{\varepsilon}{2},
\end{aligned}$$

where the inequality follows from (28). Squaring the overlap gives the following lower bound on the fidelity:

$$F(\omega^{\text{sep}}_{A:BR}, \omega_{A:BR}) \geq 1 - \varepsilon,$$

which imply by the Fuchs-van-de-Graaf inequalities in (9) that

$$\min_{\sigma_{A:BR}\in\mathcal{S}} \|\omega_{A:BR} - \sigma_{A:BR}\|_1 \leq 2\sqrt{\varepsilon}. \tag{29}$$

Thus, the circuit in Figure 4 transforms a YES instance of QSD to a YES instance of QSEP-CIRCUIT. We note that the above argument is reminiscent of similar ones from quantum information theory [Dev05].

We now prove that the circuit in Figure 4 transforms NO instances of QSD into NO instances of QSEP-CIRCUIT. In this case, we have the promise that the states ρ_0 and ρ_1 are nearly indistinguishable:

$$\|\rho_0 - \rho_1\|_1 \leq \varepsilon.$$

Due to the Fuchs-van-de-Graaf inequalities, we have the following lower bound on the fidelity:

$$F(\rho_0, \rho_1) \geq 1 - \varepsilon,$$

and Uhlmann's theorem implies the existence of a unitary operation U_R acting on the reference system of $|\psi_{\rho_1}\rangle_{RS}$ such that

$$\langle \psi_{\rho_0}|_{RS}\, U_R \otimes I_S |\psi_{\rho_1}\rangle_{RS} \geq \sqrt{1-\varepsilon}.$$

(A global phase can be fixed for U_R such that the overlap is a real number.) Thus, Bob can apply the following controlled-unitary to the state $|\varphi\rangle_{ABRS}$:

$$C^U_{BR} \equiv |0\rangle\langle 0|_B \otimes I_R + |1\rangle\langle 1|_B \otimes U_R,$$

leading to

$$(|0\rangle\langle 0|_B \otimes I_R + |1\rangle\langle 1|_B \otimes U_R)\,|\varphi\rangle_{ABRS} = \frac{1}{\sqrt{2}}\left(|0\rangle_A |0\rangle_B |\psi_{\rho_0}\rangle_{RS} + |1\rangle_A |1\rangle_B U_R \otimes I_S |\psi_{\rho_1}\rangle_{RS}\right).$$

Then the overlap between $|\Phi^+\rangle_{AB} \otimes |\psi_{\rho_0}\rangle_{RS}$ and the resulting state is large:

$$\begin{aligned}
&\frac{1}{2}\left(\left(\langle 0|_A\langle 0|_B + \langle 1|_A\langle 1|_B\right)\otimes\langle\psi_{\rho_0}|_{RS}\right)\left(|0\rangle_A|0\rangle_B|\psi_{\rho_0}\rangle_{RS} + |1\rangle_A|1\rangle_B U_R\otimes I_S|\psi_{\rho_1}\rangle_{RS}\right)\\
&= \frac{1}{2} + \frac{1}{2}\langle\psi_{\rho_0}|_{RS}\, U_R \otimes I_S\, |\psi_{\rho_1}\rangle_{RS}\\
&\geq \frac{1}{2} + \frac{1}{2}\sqrt{1-\varepsilon}\\
&\geq \sqrt{1-\varepsilon},
\end{aligned}$$

implying that the fidelity is larger than $1-\varepsilon$. Thus, by a local operation, Bob can distill a state which is $2\sqrt{\varepsilon}$-close in trace distance to the product state $|\Phi^+\rangle_{AB} \otimes |\psi_{\rho_0}\rangle_{RS}$:

$$\left\| C^U_{BR}\,|\varphi\rangle\langle\varphi|_{ABRS}\,(C^U_{BR})^\dagger - |\Phi^+\rangle\langle\Phi^+|_{AB} \otimes |\psi_{\rho_0}\rangle\langle\psi_{\rho_0}|_{RS}\right\|_1 \leq 2\sqrt{\varepsilon}.$$

(We remark that the above argument is similar to a "decoupling" argument well known in quantum information theory [Dev05, ADHW09].)

Now, we would like to argue that the one-way LOCC distance between $\omega_{A:BR}$ and the separable state $\sigma^*_{A:BR} \in S$ closest to $\omega_{A:BR}$ is larger than an appropriate constant, so that we can claim that the circuit in Figure 4 maps NO instances of QSD to NO instances of QSEP-CIRCUIT. In order to do so, Bob first performs the local unitary C^U_{BR}. This transforms the state $(C^U_{BR})^\dagger\left(\Phi^+_{AB} \otimes (\psi_{\rho_0})_R\right) C^U_{BR}$ to $\Phi^+_{AB} \otimes (\psi_{\rho_0})_R$ and the separable state $\sigma^*_{A:BR}$ to some other separable state $(\sigma^*_{A:BR})'$. Alice and Bob then perform a Bell experiment, guessing the state to be $|\Phi^+\rangle_{AB}$ if there is a violation of a Bell inequality and guessing a separable state otherwise [Bel64]. Equivalently, Alice and Bob could proceed as in the CHSH game (a reformulation of a Bell experiment as a nonlocal game [CHTW04]). In such a protocol, Alice flips a coin x and chooses one of two binary-outcome measurements to perform on her qubit. She sends both x and the measurement outcome a to Bob. Bob then flips a coin with outcome y and performs one of two binary-outcome measurements on his qubit, naming the measurement result b. Bob declares the state to be the Bell state in the case that $x \wedge y = a \oplus b$ (when they "win the CHSH game") and otherwise declares that it is not the Bell state. It is well known that the winning probability of the CHSH game with a Bell state is equal to $\cos^2(\pi/8) \approx 0.85$, while the maximum probability with which they can win this game with a separable state is equal to 0.75 [CHTW04]. This gives the following lower bound on the one-way LOCC distance between $(C^U_{BR})^\dagger\left(\Phi^+_{AB} \otimes (\psi_{\rho_0})_R\right) C^U_{BR}$ and $\sigma^*_{A:BR}$:

$$\begin{aligned}
\left\|(C^U_{BR})^\dagger\left(\Phi^+_{AB}\otimes(\psi_{\rho_0})_R\right)C^U_{BR} - \sigma^*_{A:BR}\right\|_{1-\mathrm{LOCC}} &= \left\|\Phi^+_{AB}\otimes(\psi_{\rho_0})_R - (\sigma^*_{A:BR})'\right\|_{1-\mathrm{LOCC}}\\
&\geq \left\|\left(\cos^2(\pi/8), \sin^2(\pi/8)\right) - (0.75, 0.25)\right\|_1\\
&\geq 0.2. \qquad (30)
\end{aligned}$$

Thus, by combining with the distillation argument above, we have the following lower bound on the one-way LOCC distance between $\omega_{A:BR}$ and $\sigma^*_{A:BR}$:

$$\begin{aligned}
\left\|\omega_{A:BR} - \sigma^*_{A:BR}\right\|_{1-\mathrm{LOCC}} &\geq \left\| \left(C^U_{BR}\right)^\dagger \left(\Phi^+_{AB} \otimes (\psi_{\rho_0})_R\right) C^U_{BR} - \sigma^*_{A:BR} \right\|_{1-\mathrm{LOCC}} \\
&\quad - \left\| \left(C^U_{BR}\right)^\dagger \left(\Phi^+_{AB} \otimes (\psi_{\rho_0})_R\right) C^U_{BR} - \omega_{A:BR} \right\|_{1-\mathrm{LOCC}} \\
&\geq \left\|\Phi^+_{AB} \otimes (\psi_{\rho_0})_R - (\sigma^*_{A:BR})'\right\|_{1-\mathrm{LOCC}} \\
&\quad - \left\| \left(C^U_{BR}\right)^\dagger \left(\Phi^+_{AB} \otimes (\psi_{\rho_0})_R\right) C^U_{BR} - \omega_{A:BR} \right\|_1 \\
&\geq 0.2 - 2\sqrt{\varepsilon},
\end{aligned} \tag{31}$$

where the second inequality follows from (6) and the fact that $\left(C^U_{BR}\right)$ is a local unitary, and the third from (30) and the argument at the end of the previous paragraph. Thus, as long as ε is small enough (so that $0.2 - 4\sqrt{\varepsilon} > 0$), there is a gap between (29) and (31). In fact, Watrous showed that it is possible to make ε exponentially small with only polynomial overhead for any instance of QSD [Wat02] by exploiting a "quantized" version of the polarization lemma in [SV97]. Thus, any protocol for deciding QSEP-CIRCUIT could also decide QSD, implying that QSEP-CIRCUIT is QSZK-hard. ■

Ideally, we would like to show that QSEP-CIRCUIT is a complete promise problem for QIP(2), but it is not clear to us how to do so. The obvious way to attempt this would be to reduce CLOSE-IMAGE to QSEP-CIRCUIT, but the problem is that CLOSE-IMAGE requires a general channel, whereas our protocol for QSEP-CIRCUIT has a very specific channel (one that applies a random permutation to the B systems and discards the last $k-1$ of them). Alternatively, we could attempt to find a QSZK proof system for QSEP-CIRCUIT, but the protocol that we have given to show that QSEP-CIRCUIT $\in$ QIP(2) does not satisfy the zero-knowledge property because, in the case of a YES instance, the verifier ends up with a state close to a k-extension of ρ_{AB}, which he could not have generated himself using a polynomial-time quantum circuit.

6 NP-hardness of the quantum separability problem

We now prove NP-hardness of QSEP-CIRCUIT by finding a reduction to it from the NP-hard matrix version of the quantum separability problem. The essence of the reduction is Knill's efficient encoding of a density matrix description of a state ρ_{AB} as a description of a quantum circuit to generate it [Kni95]. We begin by recalling the matrix version of the quantum separability problem:

Problem 10 (WMEM$_\varepsilon(M,N)$) *Given a density matrix $\rho_{AB} \in \mathcal{D}(\mathcal{H}_M \otimes \mathcal{H}_N)$ with rational entries subject to the promise that either (i) $\rho_{AB} \in \mathcal{S}$ or (ii) $\min_{\sigma_{AB}\in\mathcal{S}} \left\|\rho_{AB} - \sigma_{AB}\right\|_2 \geq \varepsilon$, with ε no smaller than an inverse polynomial in MN, decide which is the case.*

Gharibian showed that the above promise problem is NP-hard [Gha10], so our task is just to find a reduction from WMEM$_\varepsilon(M,N)$ to QSEP-CIRCUIT. First, consider that we can diagonalize the matrix ρ_{AB} in time polynomial in $MN\log(MN/\varepsilon_1)$, where ε_1 is an error parameter characterizing the precision of the diagonalization in the trace distance. We then compute a purification $|\phi_\rho\rangle_{RAB}$ of ρ_{AB} to a reference system with dimension no larger than MN. Knill's algorithm gives a quantum circuit running on $O(\log(MN))$ qubits that generates the state $|\phi_\rho\rangle^{RAB}$ [Kni95], and this algorithm

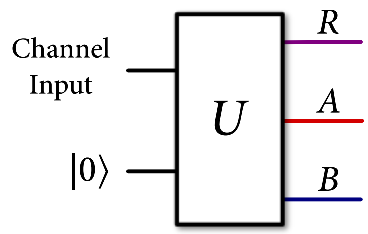

Figure 5: A quantum circuit to implement a channel. The circuit has input qubits and ancillas in the state $|0\rangle$. The circuit outputs qubits in the environment system R (which are traced over) and qubits in systems A and B.

runs in time polynomial in MN. Knill's algorithm outputs controlled single-qubit unitary gate descriptions with arbitrary precision, so we need to invoke the Solovay-Kitaev algorithm [DN06] to approximate each gate in Knill's circuit with unitaries chosen from a finite gate set, up to precision ε_2/l where l is the number of gates in Knill's circuit. The Solovay-Kitaev algorithm runs in time polylogarithmic in l/ε_2 and produces a gate sequence with length polylogarithmic in l/ε_2. This whole procedure leads to a mixed state quantum circuit generating a state ρ'_{AB} such that $\|\rho_{AB} - \rho'_{AB}\|_1 \leq \varepsilon_1 + \varepsilon_2$. The state ρ'_{AB} will be used as the input to $\mathsf{QSEP\text{-}CIRCUIT}(\delta_c, \delta_s)$.

Setting $\delta_c = \varepsilon_1 + \varepsilon_2$ implies that any instance of $\mathsf{WMEM}_\varepsilon$ for which $\rho_{AB} \in \mathcal{S}$, meaning case *(i)* of the promise, gets mapped to a YES instance of $\mathsf{QSEP\text{-}CIRCUIT}(\delta_c, \delta_s)$. For case *(ii)*, we know from Matthews *et al.* [MWW09] that

$$\min_{\sigma_{AB} \in \mathcal{S}} \|\rho_{AB} - \sigma_{AB}\|_{1-\mathrm{LOCC}} \geq \frac{1}{\sqrt{153}} \min_{\sigma_{AB} \in \mathcal{S}} \|\rho_{AB} - \sigma_{AB}\|_2 \geq \frac{\varepsilon}{\sqrt{153}}.$$

This in turn implies that

$$\min_{\sigma_{AB} \in \mathcal{S}} \|\rho'_{AB} - \sigma_{AB}\|_{1-\mathrm{LOCC}} \geq \frac{\varepsilon}{\sqrt{153}} - \varepsilon_1 - \varepsilon_2$$

so if we choose $\delta_s = \varepsilon/\sqrt{153} - \varepsilon_1 - \varepsilon_2$ then case *(ii)* gets mapped to a NO instance of $\mathsf{QSEP\text{-}CIRCUIT}(\delta_c, \delta_s)$. Moreover, because $\varepsilon_1 + \varepsilon_2$ can be made to shrink exponentially with the circuit size, the gap $\delta_s - \delta_c$ remains inverse polynomial in the circuit size. In particular, the instance of $\mathsf{QSEP\text{-}CIRCUIT}(\delta_c, \delta_s)$ will be in $\mathsf{QIP}(2)$ for sufficiently small ε, as determined by the promise in Theorem 8.

7 QIP-completeness of the channel quantum separability problem

There is a straightforward variation of $\mathsf{QSEP\text{-}CIRCUIT}$ which is a complete promise problem for $\mathsf{QIP}(3)$ (and therefore complete for QIP [KW00]). In this variation, the input is a description of a circuit that implements a quantum channel with input system S and two output systems A and B. (The channel is implemented by a unitary circuit with qubits in an environment system R that are traced over.) Figure 5 depicts a circuit that implements such a channel. The task is to decide whether there is an input to the channel such that the output state on systems A and B is separable.

Problem 11 (QSEP-CHANNEL(δ_c, δ_s)) *Given is a mixed-state quantum circuit to generate the channel $\mathcal{N}_{S\to AB}$, having an n-qubit input and an m-qubit output, along with a labeling of the qubits in the environment system R and the output qubits for A and B. Decide whether*

1. *Yes: There is an input to the channel ρ_S such that the channel output $\mathcal{N}_{S\to AB}(\rho_S)$ is δ_c-close in trace distance to a separable state $\sigma_{AB} \in \mathcal{S}$:*

$$\min_{\rho_S,\ \sigma_{AB}\in\mathcal{S}} \left\| \mathcal{N}_{S\to AB}(\rho_S) - \sigma_{AB} \right\|_1 \leq \delta_c. \tag{32}$$

2. *No: For all channel inputs ρ_S, the channel output $\mathcal{N}_{S\to AB}(\rho_S)$ is at least δ_s-far in 1-LOCC distance to a separable state:*

$$\min_{\rho_S,\ \sigma_{AB}\in\mathcal{S}} \left\| \mathcal{N}_{S\to AB}(\rho_S) - \sigma_{AB} \right\|_{1-LOCC} \geq \delta_s.$$

Theorem 12 *QSEP-CHANNEL(δ_c, δ_s) is QIP-complete if there are polynomial-time computable functions $\delta_c, \delta_s : \mathbb{N} \to [0,1]$, such that the difference $\delta_s^2/8 - 2\sqrt{\delta_c}$ is larger than an inverse polynomial in the circuit size.*

Proof. The proof of this theorem is almost identical to the proofs of Theorems 8 and 9.

We first show that there is a three-message quantum interactive proof system for QSEP-CHANNEL. This is just the obvious modification of the circuit in Figure 3 so that it becomes a three-message proof system as in Figure 2(a). In particular, the prover first prepares a state and sends it to the verifier. The verifier inputs this state to the circuit that implements the channel $\mathcal{N}_{S\to AB}$, and the rest of the proof system proceeds as in Figure 3. In the case of a positive instance, the prover can compute the states ρ_S and σ_{AB} in (32) from the description of the channel $\mathcal{N}_{S\to AB}$. He generates ρ_S with his first unitary operation and then proceeds by choosing his second unitary operation as if the state $\mathcal{N}_{S\to AB}(\rho_S)$ were σ_{AB}. Following the same analysis as in the proof of Theorem 8, the maximum probability with which the verifier accepts in this case is no smaller than $1 - 2\sqrt{\delta_c}$. In the case of a negative instance, by Lemma 1, for every state $\mathcal{N}_{S\to AB}(\rho_S)$, there is some k polynomial in the circuit size such that the maximum probability with which the prover can make the verifier accept is no larger than $1 - \delta_s^2/8$. An upper bound on the maximum acceptance probability is

$$\max_{\omega_S,\ \sigma_{AB}\in\mathcal{E}_k} F\left(\mathcal{N}_{S\to AB}(\omega_S), \sigma_{AB}\right),$$

a formula which follows from (14) and our previous analysis in the proof of Theorem 8. This leaves a gap of $\delta_s^2/8 - 2\sqrt{\delta_c}$ between completeness and soundness error (promised to be larger than an inverse polynomial) and it is known that this gap can be amplified [KW00]. Thus, QSEP-CHANNEL(δ_c, δ_s) $\in$ QIP.

To show that QSEP-CHANNEL is QIP-hard, it suffices to exhibit a reduction from the QIP-complete promise problem QCD (Problem 4) to QSEP-CHANNEL. This reduction is essentially the same as that in the proof of Theorem 9, except that the circuit in Figure 4 is modified so that the unitaries being controlled are the unitaries that generate the channels (rather than the ones that generate the states). In the case of a positive instance of QCD, there exists an input to the channels such that their outputs are nearly distinguishable, so that the output of the modified circuit is nearly separable. Also, in the case of a negative instance, the outputs of the channels for all inputs are nearly indistinguishable, so that it is possible to distill a Bell state from the output state of the modified circuit. The CHSH game argument then applies as well. Thus, QSEP-CHANNEL is QIP-hard. ■

8 A two-message quantum interactive proof system decides the multipartite quantum separability problem

After their seminal work on k-extendibility as a test of separability for bipartite states [DPS02, DPS04], Doherty *et al.* developed a notion of k-extendibility for multipartite quantum states [DPS05]. Recently, Brandão and Christandl exploited this notion to construct a quasi-polynomial time algorithm that decides a variant of the multipartite quantum separability problem [BC12]. Brandão and Harrow then improved the runtime of the algorithm by proving a stronger quantum de Finetti theorem that is applicable to the multipartite case [BH12].

In this section, we extend our results from Section 4 to the multipartite case by using the results in [BH12]. In particular, we formulate a variant of the multipartite quantum separability problem that we name MULTI-QSEP-CIRCUIT.

We begin with a few definitions before proceeding to the main theorem of this section. A multipartite quantum state with l parties $A_1, \ldots, A_l$ is fully separable if it can be written in the following form:

$$\sigma_{A_1:\cdots:A_l} = \sum_{x\in\mathcal{X}} p_X(x)\ \ |\psi_x^1\rangle\langle\psi_x^1|_{A_1} \otimes \cdots \otimes |\psi_x^l\rangle\langle\psi_x^l|_{A_l} \tag{33}$$

Let $\mathcal{S}$ denote the set of l-partite fully separable states. (We omit the dependence on the number of parties, l, which should be clear from context.)

The notion of k-extendibility extends in a natural way to multipartite systems. For notational simplicity, we refer to the total system which we are extending as C, and the l subsystems of C as $A_1, A_2, \ldots, A_l$. For example, we abbreviate $\sigma_{A_1:\cdots:A_l}$ simply as σ_C. A multipartite state $\rho_C = \rho_{A_1:\cdots:A_l} \in \mathcal{D}(\mathcal{H}_A \otimes \cdots \otimes \mathcal{H}_{A_l})$ is k-extendible [DPS05] if there exists a state $\omega_{CC_2\cdots C_k} \in \mathcal{D}(\mathcal{H}_C \otimes \mathcal{H}_{C_2} \otimes \cdots \otimes \mathcal{H}_{C_k})$ such that

1. Each Hilbert space $\mathcal{H}_{C_{i,j}}$ is isomorphic to $\mathcal{H}_{A_j}$ for all $i \in \{1, \ldots, k\}$ and $j \in \{1, \ldots, l\}$. (We are using the notation $C_{i,j}$ to refer to the j^{th} subsystem of C_i.)

2. For all parties $j \in \{1, \ldots, l\}$, the state $\omega_{CC_2\cdots C_k}$ is invariant under permutations of the systems $C_{1,j}$ through $C_{k,j}$. Note that there are $l \cdot k!$ such permutations.

3. The state $\omega_{CC_2\cdots C_k}$ is an extension of ρ_C:

$$\rho_C = \mathrm{Tr}_{C_2\cdots C_k}\{\omega_{CC_2\cdots C_k}\}.$$

Let $\mathcal{E}_k$ denote the set of k-extendible states for l parties (again suppressing the dependence of $\mathcal{E}_k$ on l as it should be clear from context). A fully separable state $\sigma_{A_1:\cdots:A_l}$ of the form in (33) has a k-extension of the following form for all k:

$$\sum_{x\in\mathcal{X}} p_X(x)\ \ \left(|\psi_x^1\rangle\langle\psi_x^1|_{A_1}\right)^{\otimes k} \otimes \cdots \otimes \left(|\psi_x^l\rangle\langle\psi_x^l|_{A_l}\right)^{\otimes k}, \tag{34}$$

which can be purified as

$$|\phi_{\sigma,k}\rangle_{R'CC_2\cdots C_k} \equiv \sum_{x\in\mathcal{X}} \sqrt{p_X(x)}|x\rangle_{R'}\ \ \left(|\psi_x^1\rangle_{A_1}\right)^{\otimes k} \otimes \cdots \otimes \left(|\psi_x^l\rangle_{A_l}\right)^{\otimes k} \tag{35}$$

Problem 13 (MULTI-QSEP-CIRCUIT(δ_c, δ_s)) *Given is a mixed-state quantum circuit to generate the n-qubit state ρ_C, along with a labeling of the qubits in the reference system R and the output qubits for each system $A_1, \ldots, A_l \in C$. Decide whether*

1. *Yes: There is a fully separable state $\sigma_C \in \mathcal{S}$ that is δ_c-close to ρ_C in trace distance:*

$$\min_{\sigma_C \in \mathcal{S}} \|\rho_C - \sigma_C\|_1 \leq \delta_c.$$

2. *No: All fully separable states are at least δ_s-far from ρ_C in 1-LOCC distance:*

$$\min_{\sigma_C \in \mathcal{S}} \|\rho_C - \sigma_C\|_{1-LOCC} \geq \delta_s.$$

The promise on negative instances is in terms of the multipartite 1-LOCC distance defined in [BH12]:

$$\|\rho_{A_1:\cdots:A_l} - \sigma_{A_1:\cdots:A_l}\|_{1-\mathrm{LOCC}} \equiv \max_{\Lambda_2,\cdots,\Lambda_l} \|(I_{A_1} \otimes \Lambda_2 \otimes \cdots \otimes \Lambda_l)(\rho_{A_1:\cdots:A_l} - \sigma_{A_1:\cdots:A_l})\|_1, \quad (36)$$

where $\Lambda_2, \cdots, \Lambda_l$ are quantum-to-classical channels.

Theorem 14 *MULTI-QSEP-CIRCUIT$(\delta_c, \delta_s) \in$ QIP(2) if there are polynomial-time computable functions $\delta_c, \delta_s : \mathbb{N} \to [0, 1]$, such that the difference $\delta_s^2/8 - 2\sqrt{\delta_c}$ is larger than an inverse polynomial in the circuit size.*

Proof. The proof system for MULTI-QSEP-CIRCUIT is similar to that of QSEP-CIRCUIT, in the sense that it amounts to a quantum computational test for multipartite k-extendibility. We exploit a generalized version of Lemma 1 that applies to multipartite states. This lemma follows from Theorem 2 of Brandão and Harrow [BH12], and we provide a proof for it in the appendix.

Lemma 15 *Let ρ_C be ε-away in one-way LOCC distance from the set of fully separable states, for some $\varepsilon > 0$:*

$$\min_{\sigma_C \in \mathcal{S}} \|\rho_C - \sigma_C\|_{\textit{1-LOCC}} \geq \varepsilon.$$

Then the state ρ_C is δ-away in trace distance from the set of k-extendible states:

$$\min_{\sigma_C \in \mathcal{E}_k} \|\rho_C - \sigma_C\|_1 \geq \delta,$$

for $\delta < \varepsilon$ and where

$$k = \left\lceil l + \frac{4l^2 \log |C|}{(\varepsilon - \delta)^2} \right\rceil.$$

Figure 6 depicts a two-message quantum interactive proof system for MULTI-QSEP-CIRCUIT. The protocol begins with the verifier preparing the state $|\psi_\rho\rangle_{RC}$, a particular purification of ρ_C, by running the quantum circuit U_ρ as given in the problem instance. The verifier transmits the reference system to the prover, who then acts on R and some ancillary qubits with a unitary P_1 that has output systems R', C_2, ..., C_k. The prover transmits systems C_2, ..., C_k to the verifier. The verifier then performs phase estimation over the symmetric group [Kit95, BBD+97] (also known as the "permutation test" [KNY08]) on the registers C, C_2, ..., C_k, using the qubits in system D

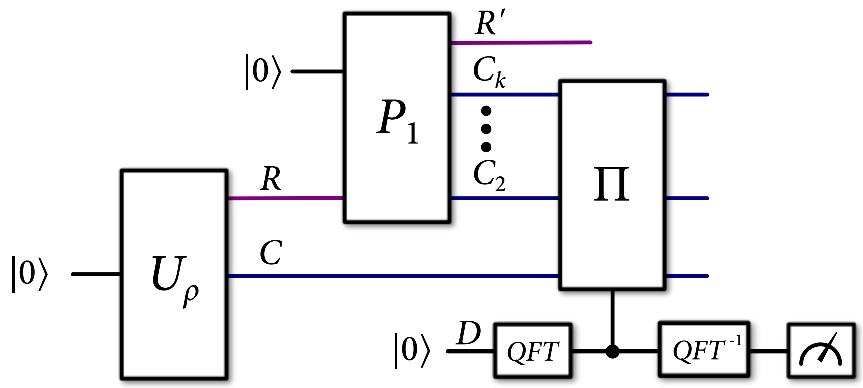

Figure 6: A two-message quantum interactive proof system for MULTI-QSEP-CIRCUIT. As in the bipartite case, the protocol begins with the verifier executing the circuit U_ρ that generates the state ρ_C. He sends the reference system to the prover. In the case that ρ_C is fully separable, the prover should be able to act with a unitary on the reference system and some ancillas in order to generate a multipartite k-extension of ρ_C to the systems C_2 through C_k. The prover sends all of the extension systems back to the verifier, who then performs phase estimation over the symmetric group (a quantum Fourier transform followed by a controlled permutation) in order to test if the state sent by the prover is a multipartite k-extension.

as the control. This control register requires $\lceil \log(l \cdot k!) \rceil$ qubits because the permutations included in the test are those from the definition of multipartite k-extendibility. The verifier performs a computational basis measurement on all of the qubits in the control register D and accepts if and only if the measurement outcome is all zeros. This protocol operates nearly exactly as the bipartite case does, except that the verifier asks for k-extensions of all the systems instead of just one. The analysis of the protocol proceeds almost identically to the analysis in Section 4, however we state it here explicitly for clarity.

In a YES instance, there is some fully separable state $\sigma_C \in \mathcal{S}$ that is δ_c-close in trace distance to ρ_C. By Uhlmann's theorem in (8) and the Fuchs-van-de-Graff inequalities in (9), there exists a purification $|\psi_\sigma\rangle$ of σ_C such that

$$\left\| |\psi_\rho\rangle\langle\psi_\rho|_{RC} - |\psi_\sigma\rangle\langle\psi_\sigma|_{RC} \right\| \le 2\sqrt{\delta_c}. \tag{37}$$

As such, the prover can operate by providing a k-extension for the separable state σ_C instead, giving a lower bound on the probability that the verifier accepts. Letting U be the unitary that includes the prover's operation P_1 and the verifier's permutation test, we have that

$$\operatorname{Tr}\left\{ |0\rangle\langle 0|_D \, U(|\psi_\rho\rangle\langle\psi_\rho|_{RC})U^\dagger \right\} \ge 1 - 2\sqrt{\delta_c},$$

where the inequality follows exactly the same line of reasoning as the steps in (17) by exploiting (37) instead.

The analysis of a NO instance also proceeds similarly to that of the bipartite case, but the channel generated by the inverse of the verifier's circuit conditional on accepting is now given by

$$\mathcal{M}_{CC_2\cdots C_k \to CD}\left(\sigma_{CC_2\cdots C_k}\right) \equiv \operatorname{Tr}_{C_2\cdots C_k}\left\{ (U_\Pi)_{CC_2\cdots C_k D}\left(\sigma_{CC_2\cdots C_k} \otimes |\mathrm{perm}\rangle\langle \mathrm{perm}|_D\right)\left(U_\Pi^\dagger\right)_{CC_2\cdots C_k D} \right\},$$

where $(U_\Pi)_{CC_2\cdots C_k D}$ is a controlled-permutation operation defined similarly to that in (19), and the permutations involved are those from the definition of multipartite k-separability. Also, $|\text{perm}\rangle$ is defined similarly to (18), though it is a uniform superposition of all possible $l \cdot k!$ permutations required from the definition of multipartite k-extendibility. Tracing out the control register D gives the following upper bound on the acceptance probability:

$$\max_{\sigma_{CC_2\cdots C_k}} F(\rho_C, \mathcal{M}_{CC_2\ldots C_k\to C}(\sigma_{CC_2\cdots C_k}), \tag{38}$$

where $\mathcal{M}_{CC_2\ldots C_k\to C}$ is a channel that applies a permutation selected at random from the multipartite k-extendibility permutations and then discards the systems $C_2, \ldots, C_k$. As in the bipartite case, since the channel $\mathcal{M}_{CC_2\cdots C_k\to C}$ symmetrizes the state of the systems $CC_2\cdots C_k$, the maximum in (38) is achieved by a state $\sigma_{CC_2\cdots C_k}$ which is permutation symmetric with respect to the multipartite k-extendibility permutations. As such, we can rewrite (38) as the maximum k-extendible fidelity of ρ_C:

$$\max_{\sigma_{CC_2\cdots C_k}} F(\rho_C, \mathcal{M}_{CC_2\ldots C_k\to C}(\sigma_{CC_2\cdots C_k}) = \max_{\sigma_C\in\mathcal{E}_k} F(\rho_C, \sigma_C). \tag{39}$$

By Lemma 15, if we take k to be larger than

$$\left\lceil l + \frac{4l^2 \log|C|}{(\delta_s - \delta_s')^2}\right\rceil,$$

then

$$\min_{\sigma_C\in\mathcal{E}_k} \|\rho_C - \sigma_C\|_1 \geq \delta_s',$$

for δ_s' strictly less than δ_s, which we can enforce by setting $\delta_s' = \delta_s/\sqrt{2}$. Observe that k is polynomial in the circuit size because the number of parties cannot exceed the number of qubits upon which the circuit acts, it is only linear in the number of qubits in system C, and the promise guarantees that $1/\delta_s^2$ is polynomial in the circuit size. Then, using the same analysis as in the bipartite case in (22)-(24), we have that

$$\max_{\sigma_C\in\mathcal{E}_k} F(\rho_C, \sigma_C) \leq 1 - \delta_s^2/8.$$

In the above we have obtained the same separation between completeness and soundness error as in the bipartite case. As discussed in Section 4, this is sufficient to place the protocol in QIP(2). (See Section 3.2 of Ref. [JUW09] for how to amplify an inverse polynomial gap.) Thus, we have given a two-message quantum interactive proof system that decides the multipartite quantum separability problem. ∎

Note that MULTI-QSEP-CIRCUIT is also QSZK- and NP-hard, as the bipartite separability problem is merely a special case of the multipartite separability problem. Also, it should be clear that if we define a "channel" variant of MULTI-QSEP-CIRCUIT that is the natural combination of MULTI-QSEP-CIRCUIT and QSEP-CHANNEL, such a promise problem is QIP-complete by the analysis in this and the previous section.

9 Conclusion

We have provided the first nontrivial example of a promise problem that is in QIP(2) and hard for both QSZK and NP. We accomplished this by introducing a version of the quantum separability

problem in which the input string is the specification of a quantum circuit that generates a mixed bipartite state ρ_{AB}, along with a promise that the state is close in trace distance to some separable state or there is no separable state close to it in 1-LOCC distance. We showed that this promise problem, called QSEP-CIRCUIT, is decidable by a two-message quantum interactive proof system, and we also proved that it is hard for quantum statistical zero knowledge (QSZK) proof systems. Our results in Section 6 also demonstrate that QSEP-CIRCUIT is hard for NP. Finally, we considered a natural variation of the quantum separability problem, in which the circuit generates a channel rather than a state, and the goal is to determine if there is an input to the circuit for which the output across some bipartite cut is separable. The "channel quantum separability problem" is complete for QIP, the class of promise problems decidable by general quantum interactive proof systems. Furthermore, we have shown that a two-message quantum interactive proof system can decide a variant of the multipartite quantum separability problem in which the input and promises are similar to those from QSEP-CIRCUIT.

Not much is currently known about two-message quantum interactive proof systems (QIP(2)), other than a nontrivial lower bound on it given by Wehner [Weh06] and the containment QSZK $\subseteq$ QIP(2) [Wat02, Wat09b]. Also, CLOSE-IMAGE is a complete promise problem for QIP(2), but it really just amounts to a trivial rewriting of the definition of QIP(2), much like the relationship between CLOSE-IMAGES and QIP(3) [KW00, RW05, Ros09]. The promise problem QSEP-CIRCUIT seems to have a natural two-message quantum interative proof system that does not satisfy the zero-knowledge property (if the verifier in our protocol were able to generate a k-extension himself, then he would be able to solve the problem efficiently in BQP, but our hardness results suggest that this should not be the case). Thus, it seems possible that one might be able to show that QSEP-CIRCUIT is in fact QIP(2)-complete, but a proof evades us for now. It is nonetheless suggestive that a simple variation of QSEP-CIRCUIT (QSEP-CHANNEL) leads to a QIP-complete promise problem. Also, since it is unclear how to begin to show that QSEP-CIRCUIT is QIP(2)-hard, it would be desirable to show that QSEP-CIRCUIT is QMA-hard, given that QMA $\subseteq$ QIP(2) and QSZK $\subseteq$ QIP(2), and QMA and QSZK are not known to be commensurate complexity classes.

10 Acknowledgements

We acknowledge useful discussions with Fernando Brandão, Anne Broadbent, Eric Chitambar, Claude Crépeau, Aram Harrow, and John Watrous. PH acknowledges support from the Canada Research Chairs program, the Perimeter Institute, CIFAR, FQRNT's INTRIQ, NSERC and ONR through grant N000140811249. The Perimeter Institute is supported by Industry Canada and Ontario's Ministry of Economic Development & Innovation. MMW acknowledges support from the Centre de Recherches Mathématiques at the University of Montreal.

A Appendix

A.1 Approximate k-extendibility

The following proposition applies to bipartite states that are approximately k-extendible:

Proposition 16 *Let ρ_{AB} be δ-close to a k-extendible state, in the sense that*

$$\min_{\sigma_{AB}\in\mathcal{E}_k} \|\rho_{AB} - \sigma_{AB}\|_1 \leq \delta, \tag{40}$$

for some $\delta > 0$, where $\mathcal{E}_k$ is the set of k-extendible states. Then, the following bound holds

$$\|\rho_{AB} - \mathcal{S}\|_{\textit{1-LOCC}} \leq \left(16\ln 2\left[\frac{\log|A|}{k}\right]\right)^{1/2} + \delta.$$

Proof. The proof of this theorem requires only a simple modification of the original proof of Brandão *et al.* [BCY11b]. First, recall that the squashed entanglement of some bipartite state ω_{AB} is equal to $E_{\text{sq}}(\omega_{A:B}) \equiv \frac{1}{2}\inf\{I(A;B|E)_{\omega'} : \omega_{AB} = \text{Tr}_E\{\omega_{ABE}\}\}$ [CW04]. Let σ'_{AB} be the state that achieves the minimum in (40), and let $\sigma'_{AB_1\cdots B_k}$ be a k-extension of σ'_{AB}. (The fact that such a minimum exists follows from the argument above (21).) Following Brandão *et al.*, observe that

$$\begin{aligned}
\log|A| &\geq E_{\text{sq}}(\sigma'_{A:B_1\cdots B_k}) \\
&\geq \sum_{i=1}^{k} E_{\text{sq}}(\sigma'_{A:B_i}) \\
&= kE_{\text{sq}}(\sigma'_{A:B}) \\
&\geq k\left[\frac{1}{16\ln 2}\min_{\sigma_{AB}\in\mathcal{S}}\left\|\sigma'_{AB} - \sigma_{AB}\right\|^2_{\text{1-LOCC}}\right].
\end{aligned}$$

The second inequality is from monogamy of squashed entanglement. The final inequality exploits a theorem of Brandão *et al.* The above implies that

$$\sqrt{16\ln 2\left(\frac{\log|A|}{k}\right)} \geq \min_{\sigma_{AB}\in\mathcal{S}}\left\|\sigma'_{AB} - \sigma_{AB}\right\|_{\text{1-LOCC}}.$$

Let σ^*_{AB} be the state achieving the minimum in $\min_{\sigma_{AB}\in\mathcal{S}}\|\sigma'_{AB} - \sigma_{AB}\|_{\text{1-LOCC}}$. We then have that

$$\begin{aligned}
\left\|\sigma'_{AB} - \sigma^*_{AB}\right\|_{\text{1-LOCC}} + \delta &\geq \left\|\sigma'_{AB} - \sigma^*_{AB}\right\|_{\text{1-LOCC}} + \left\|\sigma'_{AB} - \rho_{AB}\right\|_1 \\
&\geq \left\|\sigma'_{AB} - \sigma^*_{AB}\right\|_{\text{1-LOCC}} + \left\|\sigma'_{AB} - \rho_{AB}\right\|_{\text{1-LOCC}} \\
&\geq \left\|\rho_{AB} - \sigma^*_{AB}\right\|_{\text{1-LOCC}} \\
&\geq \min_{\sigma_{AB}\in\mathcal{S}}\left\|\rho_{AB} - \sigma_{AB}\right\|_{\text{1-LOCC}}.
\end{aligned}$$

From this, we conclude the statement of the proposition. ■

The following proposition applies to l-partite states $\rho_C = \rho_{A_1:A_2:\cdots:A_l}$ that are approximately k-extendible:

Proposition 17 *Let ρ_C be δ-close to a k-extendible state, in the sense that*

$$\min_{\sigma_C\in\mathcal{E}_k}\left\|\rho_C - \sigma_C\right\|_1 \leq \delta, \tag{41}$$

for some $\delta > 0$, where $\mathcal{E}_k$ is the set of k-extendible l-partite states. Then, the following bound holds

$$\|\rho_C - \mathcal{S}\|_{\textit{1-LOCC}} \leq \sqrt{\frac{4l^2\log|C|}{k-l}} + \delta$$

where the quantity on the left is multipartite 1-LOCC distance (defined in (36)) to the set of fully separable states.

Proof. Let σ'_C be the state that achieves the minimum in (41). Since this state is k-extendible, we have from Theorem 2 of [BH12] that

$$\min_{\sigma_C \in \mathcal{S}} \left\| \sigma'_C - \sigma_C \right\|_{1-\text{LOCC}} \leq \sqrt{\frac{4l^2 \log |C|}{k-l}}, \tag{42}$$

Let σ^*_C be the state achieving the minimum on the left in (42). From the premise of the theorem, it follows that

$$\begin{aligned} \left\| \sigma'_C - \sigma^*_C \right\|_{1-\text{LOCC}} + \delta &> \left\| \sigma'_C - \sigma^*_C \right\|_{1-\text{LOCC}} + \left\| \sigma'_C - \rho_C \right\|_1 \\ &\geq \left\| \sigma'_C - \sigma^*_C \right\|_{1-\text{LOCC}} + \left\| \sigma'_C - \rho_C \right\|_{1-\text{LOCC}} \\ &\geq \left\| \sigma^*_C - \rho_C \right\|_{1-\text{LOCC}} \\ &\geq \min_{\sigma_C \in \mathcal{S}} \left\| \sigma_C - \rho_C \right\|_{1-\text{LOCC}} . \end{aligned}$$

Thus,

$$\begin{aligned} \min_{\sigma_C \in \mathcal{S}} \left\| \sigma_C - \rho_C \right\|_{1-\text{LOCC}} &< \left\| \sigma'_C - \sigma^*_C \right\|_{1-\text{LOCC}} + \delta \\ &\leq \sqrt{\frac{4l^2 \log |C|}{k-l}} + \delta, \end{aligned}$$

which concludes the proof. ■

References

[ADHW09] Anura Abeyesinghe, Igor Devetak, Patrick Hayden, and Andreas Winter. The mother of all protocols: Restructuring quantum information's family tree. *Proceedings of the Royal Society A*, 465(2108):2537–2563, August 2009. arXiv:quant-ph/0606225.

[AKN98] Dorit Aharonov, Alexei Kitaev, and Noam Nisan. Quantum circuits with mixed states. In *Proceedings of the thirtieth annual ACM Symposium on Theory of Computing*, STOC '98, pages 20–30, New York, NY, USA, May 1998. ACM. arXiv:quant-ph/9806029.

[BACS07] Dominic W. Berry, Graeme Ahokas, Richard Cleve, and Barry C. Sanders. Efficient quantum algorithms for simulating sparse hamiltonians. *Communications in Mathematical Physics*, 270(2):359–371, 2007.

[BBC+93] Charles H. Bennett, Gilles Brassard, Claude Crépeau, Richard Jozsa, Asher Peres, and William K. Wootters. Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels. *Physical Review Letters*, 70(13):1895–1899, March 1993.

[BBD+97] Adriano Barenco, Andre Berthiaume, David Deutsch, Artur Ekert, Richard Jozsa, and Chiara Macchiavello. Stabilization of quantum computations by symmetrization. *SIAM Journal on Computing*, 26(5):1541–1557, October 1997. arXiv:quant-ph/9604028.

[BC12] Fernando G. S. L. Brandão and Matthias Christandl. Detection of multiparticle entanglement: Quantifying the search for symmetric extensions. *Physical Review Letters*, 109:160502, October 2012. arXiv:1105.5720.

[BCY11a] Fernando G. S. L. Brandão, Matthias Christandl, and Jon Yard. Faithful squashed entanglement. *Communications in Mathematical Physics*, 306:805–830, September 2011. arXiv:1010.1750.

[BCY11b] Fernando G. S. L. Brandão, Matthias Christandl, and Jon Yard. A quasipolynomial-time algorithm for the quantum separability problem. *Proceedings of ACM Symposium on Theory of Computation*, pages 343–351, June 2011. arXiv:1011.2751.

[Bel64] John Stewart Bell. On the Einstein-Podolsky-Rosen paradox. *Physics*, 1:195–200, 1964.

[BH12] Fernando G. S. L. Brandão and Aram W. Harrow. Quantum de Finetti theorems under local measurements with applications. October 2012. arXiv:1210.6367.

[BSST99] Charles H. Bennett, Peter W. Shor, John A. Smolin, and Ashish V. Thapliyal. Entanglement-assisted classical capacity of noisy quantum channels. *Physical Review Letters*, 83(15):3081–3084, October 1999.

[BSST02] Charles H. Bennett, Peter W. Shor, John A. Smolin, and Ashish V. Thapliyal. Entanglement-assisted capacity of a quantum channel and the reverse Shannon theorem. *IEEE Transactions on Information Theory*, 48:2637–2655, October 2002.

[BW92] Charles H. Bennett and Stephen J. Wiesner. Communication via one- and two-particle operators on Einstein-Podolsky-Rosen states. *Physical Review Letters*, 69(20):2881–2884, November 1992.

[CHTW04] Richard Cleve, Peter Hoyer, Ben Toner, and John Watrous. Consequences and limits of nonlocal strategies. *Proceedings of the 19th Annual IEEE Conference on Computational Complexity*, pages 236–249, June 2004. arXiv:quant-ph/0404076.

[CLMW10] Toby S. Cubitt, Debbie Leung, William Matthews, and Andreas Winter. Improving zero-error classical communication with entanglement. *Physical Review Letters*, 104:230503, June 2010. arXiv:0911.5300.

[CW04] Matthias Christandl and Andreas Winter. "Squashed entanglement": An additive entanglement measure. *Journal of Mathematical Physics*, 45:829–840, March 2004. arXiv:quant-ph/0308088.

[Dev05] Igor Devetak. The private classical capacity and quantum capacity of a quantum channel. *IEEE Transactions on Information Theory*, 51:44–55, January 2005. arXiv:quant-ph/0304127.

[DN06] Christopher M. Dawson and Michael A. Nielsen. The Solovay-Kitaev algorithm. *Quantum Information and Computation*, 6(1):81–95, January 2006. arXiv:quant-ph/0505030.

[DPS02] Andrew C. Doherty, Pablo A. Parrilo, and Federico M. Spedalieri. Distinguishing separable and entangled states. *Physical Review Letters*, 88:187904, April 2002. arXiv:quant-ph/0112007.

[DPS04] Andrew C. Doherty, Pablo A. Parrilo, and Federico M. Spedalieri. Complete family of separability criteria. *Physical Review A*, 69:022308, February 2004. arXiv:quant-ph/0308032.

[DPS05] Andrew C. Doherty, Pablo A. Parrilo, and Federico M. Spedalieri. Detecting multipartite entanglement. *Physical Review A*, 71:032333, March 2005. arXiv:quant-ph/0407143.

[EJ98] Artur Ekert and Richard Jozsa. Quantum algorithms: Entanglement enhanced information processing. *Philosophical Transactions of the Royal Society A*, 356:1769–1782, August 1998. arXiv:quant-ph/9803072.

[Eke91] Artur K. Ekert. Quantum cryptography based on Bell's theorem. *Physical Review Letters*, 67(6):661–663, August 1991.

[FvdG99] Christopher A. Fuchs and Jeroen van de Graaf. Cryptographic distinguishability measures for quantum-mechanical states. *IEEE Transactions on Information Theory*, 45:1216, May 1999. arXiv:quant-ph/9712042.

[Gha10] Sevag Gharibian. Strong NP-hardness of the quantum separability problem. *Quantum Information and Computation*, 10(3):343–360, March 2010. arXiv:0810.4507.

[Gur03] Leonid Gurvits. Classical deterministic complexity of Edmonds' problem and quantum entanglement. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 10–19, June 2003. arXiv:quant-ph/0303055.

[Hel69] Carl W. Helstrom. Quantum detection and estimation theory. *Journal of Statistical Physics*, 1:231–252, June 1969.

[Hel76] Carl W. Helstrom. *Quantum Detection and Estimation Theory.* Academic, New York, 1976.

[HHHH09] Ryszard Horodecki, Paweł Horodecki, Michał Horodecki, and Karol Horodecki. Quantum entanglement. *Reviews of Modern Physics*, 81:865–942, June 2009. arXiv:quant-ph/0702225.

[Hol72] Alexander S. Holevo. An analog of the theory of statistical decisions in noncommutative theory of probability. *Trudy Moscov Mat. Obsc.*, 26:133–149, 1972. English translation: Trans. Moscow Math Soc. 26, 133–149 (1972).

[Hor97] Pawel Horodecki. Separability criterion and inseparable mixed states with positive partial transposition. *Physics Letters A*, 232(5):333–339, August 1997. arXiv:quant-ph/9703004.

[Ioa07] Lawrence M. Ioannou. Computational complexity of the quantum separability problem. *Quantum Information and Computation*, 7(4):335–370, May 2007.

[JJUW10] Rahul Jain, Zhengfeng Ji, Sarvagya Upadhyay, and John Watrous. QIP = PSPACE. *Communications of the ACM*, 53(12):102–109, December 2010. arXiv:0905.1300.

[JUW09] Rahul Jain, Sarvagya Upadhyay, and John Watrous. Two-message quantum interactive proofs are in PSPACE. *50th Annual IEEE Symposium on Foundations of Computer Science*, pages 534–543, October 2009. arXiv:0905.1300.

[Kit95] Alexei Kitaev. Quantum measurements and the abelian stabilizer problem. November 1995. arXiv:quant-ph/9511026.

[Kni95] Emmanuel Knill. Approximation by quantum circuits. August 1995. arXiv:quant-ph/9508006.

[KNY08] Masaru Kada, Harumichi Nishimura, and Tomoyuki Yamakami. The efficiency of quantum identity testing of multiple states. *Journal of Physics A: Mathematical and Theoretical*, 41(39):395309, October 2008. arXiv:0809.2037.

[KW00] Alexei Kitaev and John Watrous. Parallelization, amplification, and exponential time simulation of quantum interactive proof systems. *Proceedings of the 32nd ACM Symposium on Theory of Computing*, pages 608–617, May 2000.

[MWW09] William Matthews, Stephanie Wehner, and Andreas Winter. Distinguishability of quantum states under restricted families of measurements with an application to quantum data hiding. *Communications in Mathematical Physics*, 291(3):813–843, November 2009. arXiv:0810.2327.

[NC00] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information.* Cambridge University Press, 2000.

[PLM+11] R. Prevedel, Y. Lu, W. Matthews, R. Kaltenbaek, and K. J. Resch. Entanglement-enhanced classical communication over a noisy classical channel. *Physical Review Letters*, 106:110505, March 2011. arXiv:1010.2566.

[Ros09] Bill Rosgen. *Computational Distinguishability of Quantum Channels.* PhD thesis, University of Waterloo, September, 2009. arXiv:0909.3930.

[RW05] Bill Rosgen and John Watrous. On the hardness of distinguishing mixed-state quantum computations. *Proceedings of the 20th IEEE Conference on Computational Complexity*, pages 344–354, June 2005. arXiv:cs/0407056.

[Sti55] W. F. Stinespring. Positive functions on C*-algebras. *Proceedings of the American Mathematical Society*, 6:211–216, 1955.

[SV97] Amit Sahai and Salil Vadhan. A complete promise problem for statistical zero-knowledge. In *Proceedings of the 38th Annual Symposium on Foundations of Computer Science*, FOCS '97, pages 448–457, Washington, DC, USA, October 1997. IEEE Computer Society.

[Uhl76] Armin Uhlmann. The "transition probability" in the state space of a *-algebra. *Reports on Mathematical Physics*, 9(2):273–279, 1976.

[VV12] Umesh Vazirani and Thomas Vidick. Fully device independent quantum key distribution. October 2012. arXiv:1210.1810.

[Wat02] John Watrous. Limits on the power of quantum statistical zero-knowledge. *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science*, pages 459–468, November 2002. arXiv:quant-ph/0202111.

[Wat03] John Watrous. PSPACE has constant-round quantum interactive proof systems. *Theoretical Computer Science*, 292(3):575–588, January 2003.

[Wat04] John Watrous. Lecture 17: LOCC distinguishability of sets of states. Theory of Quantum Information (course lecture notes), 2004.

[Wat09a] John Watrous. Quantum computational complexity. *Encyclopedia of Complexity and System Science*, 2009. arXiv:0804.3401.

[Wat09b] John Watrous. Zero-knowledge against quantum attacks. *SIAM Journal on Computing*, 39(1):25–58, 2009. arXiv:quant-ph/0511020.

[Weh06] Stephanie Wehner. Entanglement in interactive proof systems with binary answers. In Bruno Durand and Wolfgang Thomas, editors, *STACS 2006*, volume 3884 of *Lecture Notes in Computer Science*, pages 162–171. Springer Berlin / Heidelberg, February 2006. arXiv:quant-ph/0508201.

[Wer89] Reinhard F. Werner. Quantum states with Einstein-Podolsky-Rosen correlations admitting a hidden-variable model. *Physical Review A*, 40:4277–4281, October 1989.